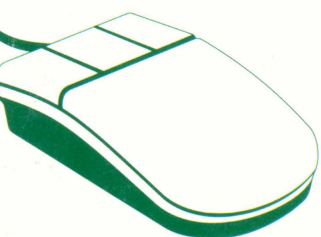


TECH TIPS

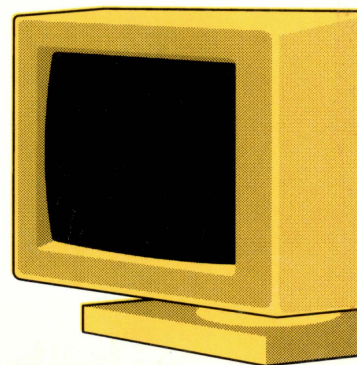
Your Computer's

DIGEST

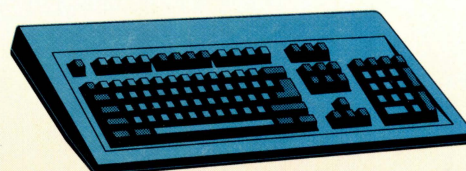
WHAT'S BEST — A SERIAL OR BUS MOUSE?



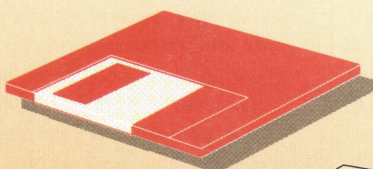
**CAN I USE MY OLD
MULTI-STANDARD
MONITOR FOR VGA?**



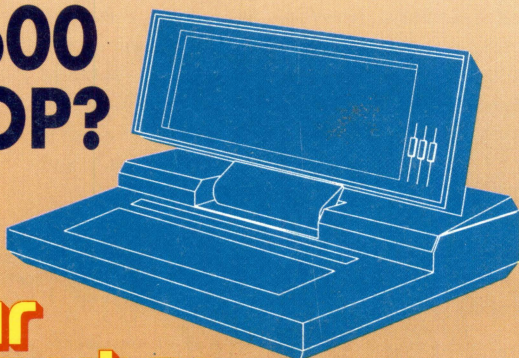
**HOW DO I PROGRAM FUNCTION
KEYS WITH ANSI.SYS?**



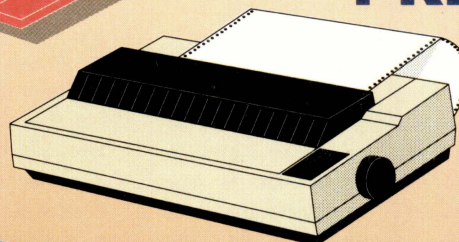
**HOW CAN I READ MAC
DISKS IN A PC 3.5-INCH
DRIVE?**



**HOW CAN I
ATTACH A
SCANNER TO
A T1600
LAPTOP?**

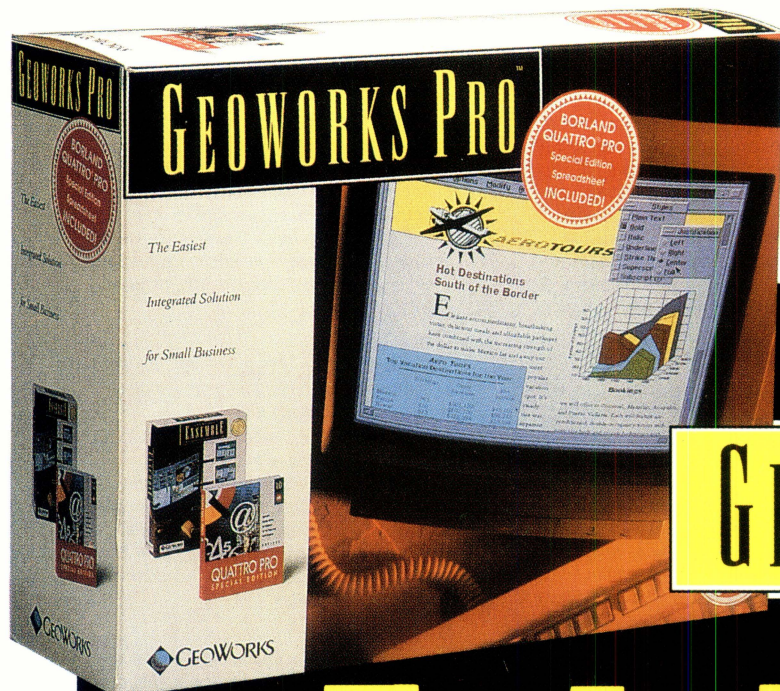


**HOW CAN I
USE A SERIAL
PRINTER ON A
PARALLEL
PORT?**



**HOW CAN I
CONVERT
CP/M FILES
TO MAC
FORMAT?**





GEOWORKS PRO™

The Easiest Integrated Solution for Small Business

GeoWorks Pro™ delivers two great products for the price of one — GeoWorks Ensemble™ and Quattro® Pro Special Edition. Together, these two award-winning products give you the easiest and most complete integrated solution for your small business computer. **Everything You Need in One Package** Word processor with spell checker, spreadsheet with charting and database, graphics, communications, address book with label maker, appointment scheduler, file and disk manager, and *much more*. **No Experience Required.** Even if you're brand new to computers, GeoWorks Pro is a breeze to use. GeoWorks Ensemble is based on an incredibly easy, windowing, graphical environment. All the applications in Ensemble work the same easy "point-and-click" way. Quattro Pro Special Edition features intuitive pull-down menus and dialogue boxes which make it the easiest to use full-featured

spreadsheet for the PC. **The Only Small Business Solution with Quattro Pro Special Edition** GeoWorks Pro is the only integrated solution for small business with a powerful, professional level, Lotus® 1-2-3® compatible Quattro Pro Special Edition spreadsheet from Borland. When it comes to managing your company's finances, you can't afford to compromise. With GeoWorks Pro you don't. **Engineered for Integration** Quattro Pro Special Edition and Ensemble work great together. The specially designed Quattro Pro Viewer lets you bring your spreadsheet numbers and graphs into Ensemble documents quickly and easily. **The Best Printing Available** GeoWorks Ensemble's advanced printing technology gives you smooth, laser quality output even on 9-pin dot-matrix printers. **Networks Supported** GeoWorks Pro lets you share files and printers on popular networks like Novell®, LANtastic™, and LAN Manager™.

ONLY \$249.00



COMPUTERMATE®
(australia) Pty. Ltd.

P.O. Box 64, Mt. Kuring-Gai, N.S.W. 2080
Ph: (02) 457 8388 Fax: (02) 457 8739

ACN 003 868 497

All rights reserved. Geoworks, Geoworks Pro, PC/GEOS, GeoManager, GeoWrite, GeoDraw, GeoPlanner, GeoComm, are trademarks of Geoworks. GeoDex is a registered trademark of Geoworks. GeoWorks Ensemble contains Nimbus Q from Digital Typeface Corp. and typefaces from URW. All other trademarks belong to companies other than GeoWorks. Actual program screens may vary slightly from screens shown.

INTRODUCTION

I had a great deal of pleasure in compiling this collection of Mark Cheeseman's *Tech Tips* culled from the pages of *Your Computer* magazine.

I particularly enjoyed, on re-reading, the little homilies, cautions, admonitions, wry remarks, asides, the humour, riding along with the sound, solid technical advice. Perhaps I should have called this *The Thoughts of Chairman Cheeseman*, or some such.

Reading these *Tips* allows us an insight into Mr. Cheeseman, as well as solving our PC problems. He is revealed to be: —

Morally upright: "And if you like it, don't forget to register!";

Whimsical: using the multi-tasking capabilities of Windows to enable a PC to play itself at Reversi really shows where he's coming from;

Blunt: "SideKick is one of the worst behaved programs ever created ...";

Plaintive: "Although we have all heard about such a card recently, we don't know who is distributing it. Can anyone out there help?" (And *Your Computer* readers being the people they are, back came the required information post haste);

Expert in matters religious: "The Mac's implementation of RS-422 is not strictly kosher, ...";

Un-Australian, (he reads the manuals): "Believe it or not, it's in the manual ... , not logically indexed, perhaps, but it's there.";

Forthright: " I ... disagree with the suggestion in that book.";

Strict, when occasion demands: " ... in a word, don't.";

Definitive: "The simple answer is 'No', it can't be done.";

A consumer advocate: " ... ask for your money back.";

Physically weak: "Not much comfort for those of us who can't carry around a 20kg battery with our laptops!";

Penny-pinching: "So the data you're trying to recover has to be worth at least as much as a floppy disk.";

Of enquiring mind: "If it was me, I would probably have a play with it, if only to satisfy my curiosity.";

Persistent: " ... I haven't quite got that one to work yet. When I do, I'll let you know.";

Ironic: "It's good to see a problem from somebody who *has* backed-up their drive, for a change.";

Patient in the face of obtuseness: " It would appear that you have misunderstood the purpose of the exercise which I described ... ";

Schoolmasterly: "Copy protection of disks is a nasty habit which we had all hoped software publishers would have grown out of by now.";

Worldly-wise: "You could try to convince either the software publisher or makers ... to provide a suitable driver (good luck!), or get a new ... ";

Comforting: " ... it's perfectly normal, and nothing to worry about.";

Forbearing: "It's probably a bit late to tell you this now, ... ";

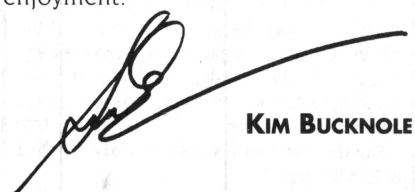
Game for anything: "Well, in lieu of a full-blown article on printers, I'll try to answer your questions here.".

ACTUALLY, Mark Cheeseman comes through as a good bloke who knows his stuff, able to impart his knowledge clearly and succinctly, as attested by several correspondents, including Clive Robertson, who finds, "... trying to make my computer work good therapy."

And we can infer much from this collection of letters about the readers and their computing: where do most problems arise? operating systems? memory management? new software? peripherals? the most widely owned CPUs? the most commonly used software? and much more.

Finally, I decided to present *Tech Tips* as they appeared chronologically in the magazine, so as not to disturb the thrust and counter-thrust of reply, follow-up, enquiry and response. Thus, when Mark refers to an earlier discussion, the reader can easily turn to it, on the fly. But given that I haven't lumped together all the memory *Tips*, all the graphics *Tips*, all the printing *Tips* and so on, I have provided other means of ready reference. In the index, key words make it possible to look up every processor poser, printer problem, Windows whinge, memory management muck-up, cross platform predicament, RS-232 riddle, modem mish-mash, and compression conundrum.

This collection of *Tech Tips* is a useful reference source, and a really good read. I hope you will share my enjoyment.



KIM BUCKNOLE

July 1990

Your Computer monthly magazine is staffed by:

EDITOR

Jake Kennedy

ASSOCIATE EDITOR

Mark Cheeseman

PRODUCTION EDITOR

Nina Stevens

ART DIRECTOR

Rolf Hagenmaier

MANAGING EDITOR

Barry Tranter

PUBLISHER

Michael Hannan

EDITORIAL AND OFFICE SERVICES

Heidi Van Es, 180 Bourke Road, Alexandria, NSW 2015. Tel: (02) 353 0959; Fax: (02) 353 0720.

ADVERTISING:**National Manager:** San Sri.

180 Bourke Road, Alexandria, NSW 2015.

Tel: (02) 353 0724; Fax: (02) 353 0720.

NSW: Gareth Owen, (02) 353 0970. **Vic:** Pilar Misa, 504, Princes Hwy, Noble Park 3174. Tel: (03) 795 3666; Fax: (03) 701 1534. **Qld:** Graham Smith, 26

Chermside St, Newstead 4006. Tel: (07) 854 1119; Fax: (07) 252 3692. **SA:** Michael Mullins, 98 Jervois St, Torrensville 5031. Tel: (08) 352 6033; Fax: (08) 352 6033. **WA:** Tony Allen, Allen & Associates, 1st Floor,

CTA Building, 54 Havelock Street, West Perth 6005. Tel: (09) 321 2998; Fax: (09) 321 2940. **New Zealand:** Federal Publishing, 67-73 View Road,

Glenfield, Auckland. Tel: (09) 443 0250; Fax: (09) 443 0249. **United Kingdom:** Peter Holloway, John Fairfax & Sons (Australia) Ltd., Associated Press House, 12 Norwich Street, London EC4A 1BH. Tel: (01) 353 9321; Fax: (01) 583 0348. **Hong Kong:**

Adrian Batten, Headway Media Services Ltd, Room 2101, Causeway Bay Centre, 15-23 Sugar Street, Hong Kong. Tel: (852) 576 8002; Fax: (852) 890 4811.

PRODUCTION MANAGER

Brett Baker

PRODUCTION CO-ORDINATOR

Patrice Wohlneck

ADVERTISING PRODUCTION

— **NSW:** Beth Parisi, (02) 353 0631;

— **Vic:** Michele Hirt, (03) 795 3666.

CIRCULATION MANAGER

Michael Prior

Distributed by Newsagents Direct Distribution Pty. Ltd., 150 Bourke Road, Alexandria, NSW 2015.

Tel: (02) 353 9911. Distributed in New Zealand by Network Distributors Ltd., 67-73 View Road,

Glenfield, Auckland. Tel: 443 0245; Fax: 443 0249.

Your Computer's Tech Tips and Your Computer are published by Federal Publishing Company, which is owned by General Newspapers Pty. Ltd., A.C.N. 000 117 322, Double Bay Newspapers Pty. Ltd., A.C.N. 000 237 598, and Brehmer Fairfax Pty. Ltd., A.C.N. 008 629 767; 180 Bourke Road, Alexandria, NSW 2015. *Your Computer's Tech Tips* is printed by Macquarie Print, Dubbo, NSW.

The Australian Publication Emblem on the front cover of this magazine signifies our pride that the editorial content in this publication is produced and wholly edited in Australia, and that most of the advertisements herein are for products and services available within Australia. © 1992

Your Computer's Tech Tips was compiled and edited by Kim Bucknole.

Monochrome and games

I own an XT with a Thomson TTL monochrome monitor and wish to run a number of games programs. As yet, I have not purchased a graphics card, as I do not know which type of card is best for my application. Can you help?

*David Fix
Mangerton, NSW*

The card which is best for you depends on the type of monochrome monitor your Thomson is. There are three major types of monochrome monitors in the PC arena at the moment — the basic MDA- and Hercules-compatible TTL unit, dual-scan TTL monitors which are compatible with both the MDA mono and CGA colour standards, and the much newer monochrome analog VGA monitors. The latter is suitable only for VGA cards, and in that application can display 64 levels of grey. This mono VGA system is the way that the industry is moving at the moment, as it means that the same VGA standard is used for both mono and color applications, and an upgrade from mono to color is performed by simply changing the monitor.

If your monitor is a dual-scan type, then an auto-switching card which supports both the CGA and Hercules standards will provide the solution. The monitor is able to detect which mode the card is currently operating in, and switch over as the card changes mode. Colors from the color card are displayed as shades of grey (or green or amber), and some foreground /background color combinations may be illegible. Since most early laptops emulated the CGA, a lot of software can be configured to provide a readable display on a monochrome screen.

If the monitor in question does

not have dual-scan capability, then there are two options. One is to use a straight Hercules-compatible card, with some form of software emulation for programs which does not support the Hercules standard, such as some games. Problems will arise, however, with software which writes directly to the screen, or games which need to be booted rather than run from the Dos prompt.

The other possibility is to use a card which connects to a standard MDA-compatible TTL monitor, but emulates both the Hercules and CGA cards. However, in either mode, the signal from the card's output port is compatible with the MDA monitor. Since the emulation is in hardware, it is compatible with auto-boot games and software which performs direct screen writes. Although we have all heard about such a card recently, we don't know who is distributing it. Can anyone out there help?

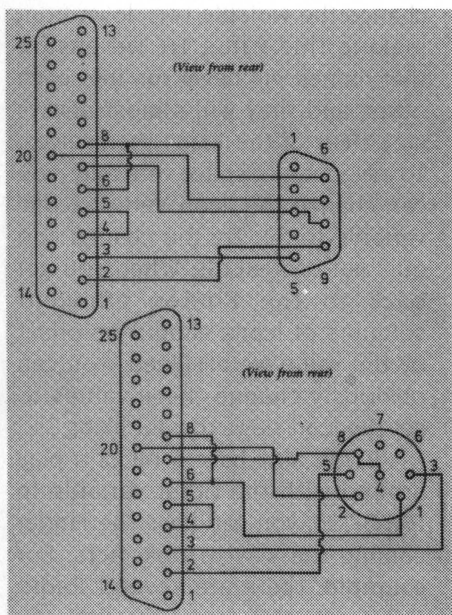
Mac modem cable

I have a Bit Blitzer 123E modem which I bought second-hand, and would like to connect it up to my Mac Plus. Is this possible, as I understand that the Mac uses the RS-422 standard for its serial ports, not RS-232?

*Keith Cowley
Zillmere, Qld.*

The Mac's implementation of RS-422 is not strictly kosher, and it is reasonably simple to construct a cable to connect your modem to the Mac. The only limitation is that, because the Mac doesn't have many handshaking lines, some of the functions cannot be used, notably Clear To Send (CTS) and Request To Send (RTS). These lines are not required for most modems,

but if you are using a modem which can have a different terminal and line speed, then problems might arise with data from the computer overrunning previously sent data. The Bit Blitzer 123E does not have different line and terminal speeds, so you shouldn't have any problem. The accompanying diagrams show how to connect a standard RS-232 device, such as a modem, to both types of Mac serial ports, either the older DB-9 or the newer 8-pin mini-din connector. (The diagrams are from Mark Cheeseman's discussion on the different types of PC connections which appeared in the January 1989 issue.)



These diagrams show 'straight-through' cables for the two types of Mac serial ports, to allow the standard RS-232C peripherals to be used with a Macintosh.

Editing TIFF files

Some time ago, I scanned some line-art and saved the images as TIFF files. I would like to import the images into PC PageMaker, but need to make some changes to the images before using in my Pagemaker documents.

*H. Richardson
Glebe, NSW*

The solution to this is reasonably roundabout, but will work provided you have the full version of Windows (not the run-time version provided with PageMaker), and don't need very high resolution. First, load the TIFF file into PageMaker (just a blank page will do), and cut or copy it to the clipboard. Then, open Windows Paint, and paste the image into it. Any minor changes can then be made to the image, which can then be either pasted back into a PageMaker document using the clipboard, or saved from Paint as a .MSP file, which can then be placed on any PageMaker page in the future, without having to re-edit the original TIFF file each time it is used.

Reading 1.44Mb disks

A colleague of mine recently sent me some data on a 3.5-inch floppy disk which I cannot read on my

machine. Since I have a high density drive, I would assume that it is capable of reading either 720K or 1.44Mb disks, but I can't read this one at all.

*Bernie Stiles,
Berourra Waters, NSW*

The problem may well be the differences between disks and disk drives. High density (HD) 3.5-inch disks differ from double density (2D) disks in two ways: the magnetic material is different; and the high density disk has a hole that indicates its type. In theory, the diskette drive can tell the difference between the two densities and will only let you format a high density disk to 1.44Mb while limiting a double density disk to 720K. In practice, many drives aren't fussy and will let you format cheap 2D disks as though they were HD. The problem comes when you try to read those disks on a drive which can tell that it's only 2D. Regardless of the fact that it can be read, the drive says 'You must be 720K but you aren't, therefore you're a bad disk!' One way to read the disk is to give it that hole it hasn't got. To find where it should be, look at the back of the disk, with the write-protect tab on the bottom left. The hole should be in the same position on the bottom right. Turn the disk over and you may find a moulding indentation where it's supposed to be. To make the hole, you can either cut it with a Stanley knife, burn it out with a soldering iron or drill it out with around a 3mm drill. Whichever you do, be

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on. The best *Tech Tip* published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is Tech Tips, Your Computer, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

careful of loose pieces or 'dags' hanging off the hold as these might damage something. You might have jumped to the obvious conclusion that this is a good way to make cheap disks do the work of expensive ones. While this will often work, the danger is that the biscuit (the magnetic 'disk') is made to a different magnetic formulation, and high density drives have different read and write electrical characteristics to double density drives. The end result is that a 'treated' disk might read correctly today, but in six months, when you really need the data, it's become flaky, and you can't get everything back.

Stretching cables

G'day fellas. When young lakey told me how he was getting some of the best PC brains to share their technical know-how with you budding troglodytes, I naturally told him I'd do my civic duty. This month I'll reveal some of the secrets of making things.

Now take cables ... how many times have you set-up the old President, only to find that the printer's six inches too far away from the PC ... or maybe it's the PC that's too far away ... Anyway, this is what you do. Just plain stretching the cable won't work. They break! The way to do it is to grow the cable. Set everything up so that the cable is nice and taught, and attach a weight somewhere near the middle. It's important to keep the PC and printer turned on while you're doing this as the cable needs sustenance to grow. All you have to do then is move the devices further apart each day, until the cable has grown long enough. Easy? Note that this process takes longer in cold rooms and works better if you

feed the cable with a spray of Mr Sheen every day.

*Vernon V. Shrunkle
JP. 970 B.ed (Retired)*

Long icon names on the Mac

I have a problem with the names of files overlapping in the Finder. While I could use shorter file names, I like the convenience of using more descriptive (and longer names). Although this is not a major problem, it would be nice to have the filenames not overlap if at all possible.

*A. Garrett
Kholo, 2ld.*

As most Mac users would know, when viewing files in icon mode in the Finder, the names of adjacent icons can often obscure each other well before reaching the Mac limit of 32-character file names. The way around this is quite easy, as long as you have ResEdit, Apple's Resource Editor — copies are available through user groups or most public domain libraries. Make sure you have a copy of all files you modify with ResEdit, because it can be harmful if you change something you shouldn't.

Once you have ResEdit, open it from the Finder. You will have to make sure you don't run ResEdit from the MultiFinder, as ResEdit won't allow you to operate on open files, and the Finder is usually open under MultiFinder. You will be presented with a list in a window similar to the 'View by Name' listing in the Finder. Just double click open the folders to take you to the Finder, then double click on the Finder itself. When you do this, you are then presented with a list of Resources within the Finder. The one we want is a 'LAYO' Resource, so double click on that listing.

Generally there should be only one LAYO resource of id #128, so once again, double click on that listing. The next listing should look a fair bit different, starting with a box for 'Font ID'. For now, we'll ignore most of the information in this resource, and concentrate on our file name dilemma.

Scroll down the listing till you come to 'Icon Horz' spacing. The value for this should be 64 (pixels, or dots on the screen). This figure sets how far apart the icons appear horizontally in the Finder. We could just assign a large number to this variable, but that would be wasteful, as icons would then only be able to go two or three across a normal window. The trick is to change the variable two down from making this value 16, we ensure that names will never run into each other, and they will almost always stay clear of the icons themselves. To save a bit of space vertically, you can change the 'Icon Vert spacing' variable to 50. Try the change and quit out of ResEdit. When you get back to the Finder, adjacent columns of icons should be offset by the value we assigned. You can obviously change these values to get the Icon views to your liking.

If you also have a bit of a play around with the other variable in LAYO 128, you can have the Finder change the way it looks/acts. For example, there are several Radio Buttons further down the list, one of which is 'Always grid drags'. By checking the '1' option, whenever you move an icon (normal or small) it will always go into a neat grid position, that is the positions the 'Clean up' option uses. This means you will almost never have to do a 'Clean up Window' again! This can also be achieved by holding down the Command Key while dragging.

A quick note: Always try to understand the changes you make. Some of the options in LAYO 128 are a bit obscure, and you could do something you don't want to. One

of the options to be particularly wary of is 'Skip trash warnings'. If you check this option, the Finder will never ask you if you are sure you want to delete something, (it asks this question for system files and applications).

Obviously, if you use this option, you will have to be fairly careful when you are deleting. Either have an undeleting utility (just in case), or if you are going to delete applications or system files, press the option key while dragging into the trash. This skips the trash warning for that delete only.

Where's the printer?

I have a problem with the parallel printer on my PC clone. Sometimes when I turn the printer on, the computer seems to think it's not there. I don't think it is a problem with the cable, as I can solve the problem by re-booting the computer. Once I have done this, the printer works fine for the rest of the computing session. Can you offer any suggestions?

*J. Everett
Surrey Hills, Vic.*

The difficulty you have is not uncommon, and relates to the way in which the BIOS searches for the presence of parallel ports in the machine during boot up. There are three locations in the I/O map that parallel ports can be located, 3BCH, 3E8H, and 2E8H. The BIOS routine checks these locations, starting at the highest one, and working its way down. Each time it finds a parallel port, it stores the location of the port in a specified area in RAM, called the BIOS data area. Dos then uses this information to assign the logical Dos names LPT1, LPT2 and LPT3, to the first, second and third ports that are found.

The way that each port is checked is to write some test data to the data port, and then perform a read operation on the same port. The ports are configured so that reading the data port allows the software to see if the data which was written to the port is what is actually being presented to the printer. If one or more of the data lines are shorted out, then the self test will not see a valid parallel port at that location, because the data written to, and read from the port do not match.

If an unpowered printer is connected to the port, the protection diodes built into the printer's input circuitry can virtually short these lines to ground. This means that the port to which the printer is connected fails the self-test, and Dos does not recognise it. If there is another parallel port in the system, then it will become LPT1, otherwise there will be no parallel ports. This can be verified by running Norton's SysInfo program (one of the Norton Utilities), and looking at the number of parallel ports that it finds — there will be one less than there should be.

If you don't have the Norton Utilities, debug can be used to look at the BIOS data area, to check the locations of the parallel ports. Enter debug, and type 'd 0040:0000', to display the start of the BIOS data area. The six bytes starting with 0040:0008 represent the I/O locations of the three parallel ports which can exist in a system, although the actual locations are not important here. A '00 00' value means that that port does not exist. If all values are zero, then there are no parallel ports in the system. Make a note of these values, and then re-boot the computer with the printer turned on, and repeat the exercise. A change in any of these values means you've found the problem.

The solution is pretty simple — always power the printer up

before the computer, or plug the two into a power board, and turn them on together.

Caps to lower case

Sometime ago I was shown how to change upper case letters to lower case in WordStar 4. I've now forgotten and have searched the Menus repeatedly for help and it's driving me mad. How is it done?

*Vera Tate
Adelaide, SA*

Easy: mark the block of text you want to change (Ctrl-KB and Ctrl-KK) and then Ctrl-K" (double quotes) will change lower case to upper; Ctrl-K' (single quote) will change upper to lower. We discovered this by accident some time ago, but it is in the manual under Case Conversion.

Trouble with Windows

I have just obtained a copy of Windows 3.0, and while I think the new user interface is very nice, I am having trouble with some of my existing applications. For example, PageMaker refuses to open a document because of a lack of memory. I am using a '386-based machine with 4Mb of memory.

*Geoff Taylor
Tullich, Vic.*

Unfortunately, there is no real solution to this which allows existing applications to access the improved memory management features of Windows 3. The immediate solution is to run Windows in real mode (by typing win /r at the Dos prompt). Unfortunately, this

disables the memory management features of the '286 standard mode, and '386 enhanced modes, but should allow existing applications to run. The best long-term solution is to obtain Windows 3-compatible versions of your application software. Microsoft assures us that they have been in contact with application developers, so that updated versions of their software was to be available soon after the launch of the new Windows.

Hint

Did you know that you can generate any extended ASCII code from the keyboard, by holding down the Alt key, and typing in the ASCII code (in decimal), on the numeric keypad? The state of the NumLock key is unimportant. This works with virtually all software (even Windows applications), although the way in which a

particular program responds to unusual keystrokes is up to the programmer. This facility is useful for drawing lines and boxes in word processing files, for example, and while some software package manuals mention it, a lot don't. The best thing to do is to try it with your applications, to see if it can be used for anything useful.

August 1990

Garbage from modem

I have a problem with my modem producing garbage. The problem does not occur all of the time, but after 12 or 13 minutes online, my screen fills with rubbish. Telecom couldn't offer any help — they said the line is OK, and the modem seems all right, since it works fine on a friend's phone line. Can you help?

*Dr David Spiniker
Fashion Village, NSW*

It sounds like you have one of Telecom's new T200 Touchfones. These phones periodically draw current from the phone line to keep their memories powered. When this occurs, the extra loading on the line causes distortion, which corrupts data communications. A quick fix is to unplug the T200 when you are using the modem — hopefully, it won't forget its memory numbers. Another alternative is to ask Telecom for a current sharing device, which allows the phone to re-charge itself without messing up the modem.

Another problem to watch is related to Easycall, specifically the Call Waiting facility. When somebody tries to call you while the modem is using the line, the call is momentarily interrupted while the

call waiting pips are transmitted. This causes garbage to be received by the modem until it determines that there is no carrier and hangs up in disgust. The simple solution is to disable the call waiting before using the modem, and to re-enable it afterwards.

Another solution is to add two entries to your dialing directory — one to disable call waiting and the other to enable it again. Then, dial the first entry at the beginning of a session, and the other after you have finished using the modem. A more automated approach might be to put the disable sequence in the modem initialisation string itself, so that the call waiting is automatically disabled whenever you fire up the comms package. This still leaves the problem of re-enabling the facility when you have finished with the modem. If your comms program can send a string to the modem before exiting, then this could be used, otherwise, you may have to resort to the dialing directory solution above for this operation.

through COM4. However, while two of the ports work without any problems (COM1 and COM3) the other two do not work reliably. I have two different modems, on COM1 and COM4, a Mouse Systems mouse on COM2, and I use COM3 to connect to my laptop with LapLink III. Whenever I try to use the modem on COM4, it does not work unless I move the mouse a bit. After that, I don't have any problems with this port, but when I exit the communications program (Telix 3.12), the mouse refuses to work again until I reboot. Do you have any suggestions?

*Jennifer Cox
Ajana, WA*

This problem is related to the way in which the communications capabilities of the PC have evolved. The original PC, and Dos, only supported two serial ports — COM1 and COM2. Each of these was allocated a specific I/O address and interrupt line. When users started demanding more I/O capabilities from their machines — especially with a mouse taking up one of the serial ports — COM3 and COM4 were introduced in some software, and in Dos itself from version 3.30 onwards.

However, because of the shortage of interrupt lines (especially on XT-level machines), they have to share the same interrupt lines as the first two. That is, COM1 and

Four serial ports

I have four serial ports on my computer (a 12MHz Everex AT), addressed as COM1

AUG 1990

COM3 share interrupt 4, and the other two share interrupt 3. Normally, this should not cause a problem since the chip which controls each serial port (the UART) has a register which the software can read to determine whether that chip caused the interrupt or not.

Problems can arise, however, when the software that is controlling the port (such as the Mouse Systems mouse driver) assumes that there is only one port on that interrupt line. It receives an interrupt from the other port on that interrupt line, but doesn't know what it is for, so it appears to hang.

One solution is to arrange the serial ports so that the shared interrupt lines do not cause conflicts.

For example, put your two modems on COM1 and COM3. The assumption here is that only one modem is to be used at once, and if you have software which can use two modems at the same time, it is smart enough to share the interrupt line between the two. This leaves COM2 for the mouse (which usually must be on COM1 or COM2), and COM4 for LapLink. Since LapLink doesn't use the interrupt line at all, there should be no conflict here.

Archiving programs

What is the best file compression/archiving software to use? I have seen ARC, ZIP and LZH file extensions, and I understand that they all use different techniques to compress data? Are all of them incompatible with one another?

*C. Williams
Jeeralang, Vic.*

As is the case with most decisions in life, there is no simple answer to this, owing to the rather checkered history of PC archiving programs in

general. The first PC archiving package to appear on the scene was Systems Enhancement Associates' ARC. This not only provided a means of compressing files to make them smaller on BBS's hard disks and faster to transmit, but also allowed several files to be combined into a single 'archive', so that a program file, its overlays and documentation could be distributed as a single unit without the danger of a vital part of the program going astray.

Since the appearance of SEA's ARC, archiving files on bulletin boards became almost universal. PKWare, an American company, thought that it could do a better job and released PKARC, which was faster than ARC in compression and decompression. However, in an effort to standardise on the ARC file format, PKARC used the same file compression techniques (with one addition) and file formats, so that files created by the two programs were compatible.

Since PKARC was faster than ARC, and had an additional compression algorithm built-in, it became the preferred compression program. SEA didn't like this, and in typical US style, took PKWare to court, although the two ultimately settled out of court. The settlement basically meant that PKWare could no longer market products which used SEA's compression techniques, and it could no longer use the term ARC. Yet another standard bites the dust!

PKWare changed PKARC and PKXARC to PKPAK and PKUNPAK, using the PAK file extension, and later changed to PKZIP and PKUNZIP. The latter programs were even faster than any of their predecessors — for a full report, have a look at John Hepworth's 'IBM Underground' column in the May 1989 issue.

A relative newcomer is LHARC, which creates files with the extension LZH. This is a favourite among

BBS Sysops because it creates the smallest files of any compression program, thus maximising valuable disk space and reducing file transfer times to an absolute minimum. The penalty is speed — it's much slower than ZIP or PKARC.

So the real answer to your question is 'it depends'. If you are compressing files for transmission over expensive phone links, or to cram the absolute maximum amount of data on a disk, then LHARC is probably for you. However, if you are looking for a quick way to backup your work directory at the end of each day before going home, then speed is probably more important, so PKZIP would be a better choice.

By the way, if you are fortunate enough to have a modem with MNP-5 data compression, don't expect it to transmit compressed files any faster than normal — MNP-5 is just another file compression system, and files which are already compressed don't compress well a second time, so the transmission speed will not improve much, if at all. Note that the slight increase experienced with MNP-4 will still occur with compressed files, since this is achieved by removing the start and stop bits before transmission, not by compressing the file.

Fast format for Dos Floppies

I don't think anybody actually likes formatting floppy disks —but unless you are prepared to squander big handfuls of dollars on factory pre-formatted ones, the occasional formatting session is a necessary task. You can only con the office junior into doing it once! A reader has come up with a simple fast formatting utility, which we have presented here for the edification of all.

Formatting is inescapable for

AUG 1990

Tech Tip of the Month — A RAM disk with Windows

WHILE THE USE of a RAM disk with Windows is not new, I am sure my hint is. In fact I'm amazed that a RAM disk with Windows has received very little press coverage, considering the significant increase in speed, at least with Corel Draw, not to mention the saving on hard drive wear and tear.

The method I use to considerably speed up Windows applications such as Corel Draw, is to create a RAM disk for the temporary files. In fact, this improvement is so efficient that Corel Draw only calls the hard drive to load a new file not already resident in Smartdrv.sys.

However, sometimes the temporary files become too large, and I run out of disk space. When that happens, it is necessary to restart Windows with the temporary files diverted to the hard disk. Care must be taken to keep a lid on the size and number of temporary files open. Keep a check on those temporary files by occasionally calling up the directory of the chosen path.

Normally, you would expect that the temporary file path could be altered at the DOS prompt, using the SET command. However, because of the possibility of losing temporary files already opened, Windows re-sets the environment according to the autoexec.bat file, so any changes are negated.

I normally use Windows with TEMP set to the RAM disk, and run it from a special batch file if I know I will need more temporary space than is available on the RAM disk.

To set up a RAM disk, include the following instructions in the nominated files. Remember, when using extended memory, that ramdrive.sys, smartdrv.sys, and himem.sys (which uses 64Kb of extended memory), should not exceed the total extended memory available.

In the autoexec.bat file, add the line

```
SET TEMP=D:\
```

Make a copy of the autoexec.bat file, called autoexec.hdd, and change the above line in this file to:

```
SET TEMP=C:\WINDOWS\TEMP
```

Finally, create the following batch file, to run Windows with the temporary files directed to the hard drive. The last line is optional.

```
@ECHO OFF
SETTEMP=C:\WINDOWS\TEMP
RENAME C:\AUTOEXEC.BAT C:\AUTOEXEC.TMP
RENAME C:\AUTOEXEC.HDD AUTOEXEC.BAT
WIN
RENAME C:\AUTOEXEC.BAT C:\AUTOEXEC.HDD
RENAME C:\AUTOEXEC.TMP C:\AUTOEXEC.BAT
SET TEMP=C:\
DOSHELL
```

PE Bailey

RAM disks are a great way of speeding up the operation of programs, if you know which files are being accessed frequently. A cache achieves pretty much the same effect, but it tries to determine which files to cache on the fly, and it often ends up caching files which aren't being accessed very often. By using a RAM disk, you can determine which files to put in the RAM disk, and which ones can stay on the hard disk. The only catch is, you have to watch very carefully the amount of space available — a RAM disk can't automatically overflow onto the hard disk if it fills up, whereas a cache can.

I can't find any evidence to suggest that Windows changes the temp variable itself, but it may check the autoexec.bat file for any TEMP assignment statements as it starts up. Running SET from within Windows reveals that the variables are as they were before starting Windows, regardless of the contents of autoexec.bat. That doesn't reduce the validity of using RAM disks, of course.

message from the new controller, and a self-test of its on-board memory. This test was rather slow, but fortunately, it only appears when the machine is first powered-up, and can be interrupted by pressing a key on the keyboard.

The sample I tried out had the full complement of RAM on-board, and I have used it in this machine for about three weeks so far, and haven't encountered any compatibility problems with Windows or DESQview, or any

other software that I've tried.

What was noticeable, however, was the dramatic improvement in speed when running disk-intensive applications. I should point out that the Syncomp has 8Mb of RAM in it, of which I usually assign half to a disk cache, when running Windows. With the DriveCache installed, I removed Smartdrive from the config.sys file entirely, leaving the full 8Mb for applications.

Systems with less memory will experience an even more dramatic

improvement, especially if you have been running without a cache of any sort, to maximise program workspace.

I performed a few quick tests, comparing the speed of various Windows applications with 4Mb of system memory only, 8Mb of system memory, with 4Mb set aside as a cache, and with 4Mb of system memory and 4Mb of memory on the cache card. Without exception, the DriveCache outperformed the Windows Smartdrive cache of the

AUG 1991

```

echo and press ENTER when
    ready . . .
pause>nul
debug<fformb.scr>nul
goto end
:noparam
REM help screen
echo.
echo The FFORM command will
    fast format a previously
echo used floppy disk by nulling
    the FATs, telling DOS
echo that the whole disk is vacant
echo.
echo No files are actually erased,
    but the sectors
echo are marked to be over-
    written in a similar way
echo to DEL *.* (but much faster)
echo.
echo FFORM will not work on
    new disks, giving a
echo "General Failure Error"
echo.
echo Type FFORM A: or FFORM
    B: whichever you require
echo.
:end

```

To put the whole scheme into operation, copy FFORM.BAT, FFORMA.SCR, and FFORMB.SCR to a directory covered by your Path. Now, from any directory, you can just type Fform (drive) and the disk will be ready in the twinkling of an eye — just the thing when you can't wait to get your latest masterpiece safe on plastic!

Seriously, I first wrote this program when I had a large number of disks to reformat, and this is the situation where it shines — it takes longer to change the disks than to fast format them.

Now if you're a real productivity whiz, you can conscript the services of the aforementioned kid to type Fform a: and Fform b: at the appropriate spots while you juggle disks in and out of both disk drives.

*J. Boetje
Broosweena, Zld.*

IDE Drives

I am considering upgrading the hard disk in my 12MHz AT. While I know all about MFM, RLL, ESDI and SCSI, I am confused by these new drives on the market which are referred to as 'AT-bus' or similar. Are these like hard cards, which plug directly into the expansion bus of the computer?

*N. Wright
Abbey, WA*

No, they're not hard cards, at least not physically. However, they are similar in that they connect directly to the expansion bus of the computer. But, they are physically mounted in a drive bay, as for any other drive, for mechanical stability, and connected by a 40-way cable. They are similar in concept to ESDI and SCSI drives in that they have the controller mounted on the drive. While SCSI and ESDI drives only have part of the controller built into the drive (and thus still require a controller card), AT drives have the entire controller built into the drive, and for that reason are also commonly known as integrated drive electronics or IDE drives.

Many computers have an IDE connector on the motherboard, or an interface card can be supplied with the drive. The IDE interface is basically an extension of the relevant lines of the AT's I/O bus. This means that the drive doesn't have to conform to any of the existing drive-interface standards since the drive and controller are one unit.

If you want to use an IDE drive on a particular machine, check that the BIOS supports the IDE drives. If it is an old BIOS, it probably won't. Ask your dealer about an upgrade to a version which does support IDE drives, or has a user-definable drive selection (recent AMI and Phoenix BIOS' have this feature). This allows the specific parameters

— number of cylinders, heads, and sectors, and the write-precompensation value — to be entered for drives which are not supported by any of the pre-defined selections.

User-definable drive types avoid the need for driver software, such as Disk Manager or SpeedStor. If you can't upgrade the BIOS, try getting one of these programs from the dealer whom you bought the drive from. This is not as elegant a solution as upgrading the BIOS, but is easier to implement (especially if you don't feel like pulling chips out).

Filename recognition

Greetings fellow PC people. Here's another technical tip from the keyboard that never sleeps. How many times have you had trouble finding a file on the old hard disk, eh? So many files and so little time. Well, uncle Vern has the answer — sensible file naming

As you already know, files have two parts to their names, technically known as 'the bit on the left of the full stop' and 'the bit on the right of the full stop'. Let's call them TBOTLOTFS and TBOTROTFS like IBM and Microsoft do. (Actually, let's call them LT for left thingy and RT for right thingy.)

Now, because the main reason for losing files is that they have different names, this is my solution — name all related files with the same LT and use different RTs to differentiate between them. So, all dBase files would be named like DBASE.NAM (my name and address book), DBASE.PHN (my phone call log), and DBASE.BTP (bills to be paid) and so on. I'm sure you'll agree this is much more sensible. Not only that, but you end up with so few extra files like indexes and backups, as you only

AUG 1990

ever need the one such as DBASE.NDX or DBASE.BAK. A great way to save space. Likewise, my wordprocessor files are all

called WORDPRSR.??? and my Lotus 1-2-3 are called ... yep, you guessed it — LOTUS123.???, and so on. So, when I sort my directory, all

my similar files are listed together. Yes, I know — brilliant.

*Vern Shrunkle
West Wobbalong*

ANSI.SYS

THE DOS ansi.sys file is a source of mystery for many people, who either install it in their config.sys file and then forget about it, or just forget about it altogether. However, with a bit of effort, it can be used to customise the screen and keyboard of the computer.

The ansi.sys driver replaces the standard console (CON:) driver in Dos, and allows much more control of the screen and keyboard than the standard driver. This control is exercised through escape codes sent to the driver, which is as simple as sending the character string to the screen with an *echo* statement called from a batch file.

The ANSI escape codes interpreted by ansi.sys are all of the form <ESC>[, followed by one or more numbers separated by semicolons, and a letter, which is the actual command. The most common command is 'm', which is used to change the character colours and attributes. For example, <ESC>[32;40;1;5m will give you flashing bright green text on a black background. You can also put these codes in your prompt statement in autoexec bat, to control the colour of the prompt. In this case, the escape code is replaced by \$E.

These codes are well documented in the Dos manual, in the appendix covering installable device drivers, so we won't present a full coverage here. ANSI escape codes can also be used to change the video mode, clear the screen,

and move the cursor around the screen. And, they can be used to re-define the function of any key on the keyboard. This latter feature is not documented at all in most Dos manuals, but is potentially one of the most useful.

For example, you can re-define the function keys, or define functions for those which are undefined. Say, you might want Alt-F10 to park the hard disk. To do this, the escape code is <ESC>[0;113;"park";13p.

The general form of a key-definition escape code is: <ESC>[N1;N2;...Nxp

The first one or two digits represent the key being defined. If the key is a normal key which has an ASCII code (such as an alphanumeric key, either alone, shifted, or with control), then the ASCII code is the one which is used here. For function keys, cursor-control keys, and Alt-key combinations, the first number is zero, followed by the code given in the table below.

All subsequent digits are the ASCII codes of the characters in the key definitions, separated by semicolons. However, a string of alphanumeric characters can be represented by putting the string itself in quotes (as in the 'park' example, above), rather than as a series of semicolon-separated ASCII codes. The '13' at the end of that particular string is of course the ASCII code for the carriage return at the end of the park command.

Alt-key combinations:

Alt- Key	Code	Alt- Key	Code	Alt- Key	Code
1	120	A	30	N	49
2	121	B	48	O	24
3	122	C	46	P	25
4	123	D	32	Q	16
5	124	E	18	R	19
6	125	F	33	S	31
7	126	G	34	T	20
8	127	H	35	U	22
9	128	I	23	V	47
0	129	J	36	W	17
-	130	K	37	X	45
=	131	L	38	Y	21
		M	50	Z	44

Note:

All key codes above are preceded by a '0'. Keys not shown here return their assigned ASCII code.

Function key codes

Key	Alone	Shift	Ctrl	Alt
F1	59	84	94	104
F2	60	85	95	105
F3	61	86	96	106
F4	62	87	97	107
F5	63	88	98	108
F6	64	89	99	109
F7	65	90	100	110
F8	66	91	101	111
F9	67	92	102	112
F10	68	93	103	113

September 1990

Monochrome card update

In 'Tech Tips' in the July issue, on the subject of dual mode cards, we mentioned a card which allows both Hercules and CGA modes to operate on a standard monochrome TTL screen, but were not sure of the supplier. One of our readers — Linda McGarry of Kentucky, NSW — tells us that the card is available from Electronic Solutions, PO Box 426, Gladesville 2111 NSW, and is called the Magic Combo Mono/Color Graphics Card. Thanks for the info, Linda.

TIFFs ain't TIFFs

Daniel Ford, from Advanced Solutions, was prompted to prepare this technical note about TIFF files when he found that the files produced by his FaxScan software (reviewed in our 'Add-on Atlas' in June), were not compatible with certain software products which can import TIFF files. It seems some TIFF implementations comply with the standard to read some TIFF files, but don't go all the way.

This note was originally prepared by Advanced Solutions for the benefit of FaxScan users. It aims to help prevent them from getting their fingers burned buying graphics software which doesn't work with FaxScan's TIFF files. However, the problem is not restricted to this program alone, so take care when you see the words 'TIFF compatible'.

TIFF is an acronym for Tagged Image File Format. It is a graphics file format administered jointly by Aldus and Microsoft, and used widely in both MS-DOS and Macintosh environments. Although not always apparent to users, TIFF files do not conform to a single, rigid 'format' (layout), but are

allowed considerable flexibility by the TIFF specification.

The current TIFF specification (TIFF5), attempts to channel some of this flexibility into more restrictive paths, to simplify the jobs of TIFF 'writers' and 'readers', (that is, programs which write or read TIFF files). It does this by defining four TIFF classes, (more may be added in some future revision).

They are class B, for bilevel (1-bit) images, class G, for grey scale images, class P for palette color images, and class R for RGB full color images. For brevity, these classes are usually referred to as TIFF B, TIFF G, TIFF P, and TIFF R. FaxScan writes only TIFF B images, so our discussion henceforth concentrates on that class.

A TIFF B file consists of four types of data — a header, an image file directory (IFD), value fields, and the actual image (pixel) data. The header, which is a fixed length of eight bytes, is always located at the beginning of the file. It simply identifies the file as a TIFF file, notes the byte order in which the data are stored, and contains the offset (location within the file), of the IFD. The IFD contains several 'tags' (from whence the name), which provides information about the type, size and resolution of the image contained in the file, as well as where the image data can be found. The tags are all of fixed length (12 bytes), and so sometimes the information they provide cannot be contained within the 12-byte tag itself. In this case, the tag data points to a Value Field, where the required information is found.

The actual image data (pixel information) is stored in image 'strips'. It is allowable, though not recommended, for an entire (large) image to be stored in a single strip. The TIFF5 specification recommends that the strip size (number of scan lines per strip), be chosen such that each strip occupies about 8K of storage, to make it

easier to buffer the image data.

The image data within the strips may be stored uncompressed or compressed. And, if it is compressed, it may be done so by one of two different schemes, (there is, of course, a tag which specifies the compression method). The compression method recommended by the TIFF5 specification is called PackBits, a simple, byte-oriented, run-length encoding scheme. The other method is Modified Huffman (MH) Coding, a complex, variable-length, bit-oriented, run-length coding scheme, similar to that used in CCITT Group 3 facsimile machines.

It is not a requirement of the TIFF specification that all TIFF B readers be able to read compressed files, though this is recommended. Many applications will refuse to read compressed TIFF files, which is a shame, because significant disk space savings can be obtained using one of the compressed formats, (depending on the image content, file sizes as small as 10 or 20 per cent of the uncompressed size can be achieved).

One central requirement of the TIFF specification is that, apart from the 8-byte header at the start of the file, all other information can be stored in any sequence anywhere within the file. A pointer (file offset), in the header points to the IFD, pointers in the IFD point to Value Fields, and pointers in the Value Fields (or sometimes in the IFD), point to the image data.

Many (including well-known), application programs, claiming to read TIFF files, ignore this most basic requirement of the TIFF specification, and assume that the various file sections will be in some specific order. These programs are noted in the Compliance Table below with an 'x' in the 'does not comply' column. Some of these, though, will read FaxScan's TIFF files, simply because we happen to have stored our TIFF data in the

0611
SEPT 1990

sequence they expect it. But they cannot be relied upon to import TIFF files from a variety of other applications, where the data might be (quite legitimately), stored in some different sequence.

FaxScan's TIFF files conform to the TIFF4 specification, simply because so many useful application programs have not yet caught up with the current (TIFF5), specification, (any many more are still at pre-TIFF4 stages of development). The TIFF4 standard was around for over a year prior to TIFF5, so there is no excuse for applications not conforming at least to that standard, especially considering that the TIFF specification is available free (to software developers), from Aldus or Microsoft.

The table below was derived from testing each of the named

programs with a set of files conforming to the TIFF4 and TIFF5 standards, and with different arrangements (sequences) of the various data components. One other factor which might affect an application's ability to handle your TIFF files, and which does not appear in the Compliance Table, is the question of image size. Whether a TIFF file is compressed or uncompressed, it must ultimately be uncompressed by the application in order to display and manipulate the image. Bitmap (for example, TIFF), images can be quite large in terms of the number of pixels. At a scanning resolution of approximately 200 x 200dpi (dots per inch), as recommended with FaxScan, an A4 page occupies almost 480K of memory, while at 200 x 400dpi an A4 page is about

960K! This exceeds the 'conventional' memory size of MS-Dos PCs, and even the 480K file cannot be entirely fitted in conventional memory if the application program (and Dos and anything else loaded in memory) occupies 160K or more.

Obviously, EMS is desirable, together with an application which knows how to use it, or else some clever programming tricks are needed to overcome the memory limitation. Designer, Optiks, PageMaker, and Ventura all handle large bitmaps (with EMS), while SLEd performs trickery, using the hard disk for 'virtual memory'. Carets/Deluxe handles a 480K page, though it is not known whether it requires EMS to do this.

Designer, the Image-In suite, and PageMaker all require Microsoft Windows to run.

Applic. Type	Application	complies		doesn't comply	reads FaxScan
		TIFF5	TIFF4		
Design	Designer (Micrografx)	✓	✓		✓
DTP	Pagemaker (Aldus)	✓	✓		✓
	Ventura Publisher (Xerox)	✓	✓		✓
OCR	Carets/By Hand Deluxe		✓		✓
	CAT Reader (Comp. Aided Tech)			x	
	Image-Read (Image-In)			x	✓
	Prodigy OCR			x	
	ReadRight OCR (OCR Systems)			x	
	Spot OCR (Flagstaff Eng'g)			x	
Paint (Pixel editors)	Astral Picture Publisher			x	
	Gray-FX (Xerox)			x	
	Image-Scan/Paint (Image-In)			x	✓
	PaintShow (Logitech)			x	✓
	Paintbrush IV for Windows			x	✓
	Paintbrush IV Plus (Zsoft)			x	✓
	SLEd (VS Software)	✓	✓		✓
Utility	Hijaak (INSET Systems)			x	
	Optiks (Graham Systems)	✓	✓		✓

SEPT 1990

Recommendations

IF YOU don't already have all the applications you need, what should you buy? The first thing to note is that we were only interested in testing the listed programs for TIFF (import), compatibility, to help users avoid obviously inappropriate choices. We did not endeavour to check other features of the programs, as this sort of subjective evaluation would naturally depend on what features a particular user desired.

The second point is that we were primarily interested in testing 'paint' and OCR programs, since a paint program may be required to 'touch up' images prior to using in a DTP application, and OCR is an application of great interest to users of FaxScan. Designer just happens to be in the list because we own a copy. It has the most competent TIFF-import facility we have seen in any program. Its only serious contender, Corel Draw was not tested, and may well be just as TIFF-capable. The two 'big guns' in DTP (PageMaker and Ventura), are obviously of interest to most scanning candidates, and we wouldn't choose between them on your behalf. Both conform to the current TIFF specification.

We were shocked and disappointed to find that the big names in paint programs just didn't comply with one of the most basic requirements of the TIFF specification. This may be because they have grown out of the 'hobby end' of the PC market, and are perhaps still written by 'hobby programmers', whose programming style is more empirical than 'by the book'.

SLEd looks like being the big winner, being the only paint program to conform to the recent TIFF standards, and having generally good editing features as well. It also exports both TIFF and PCX formats, and performs some neat tricks for printing graphics on laser

printers. But, if you don't care about global TIFF compatibility, and you intend to try OCR now or in the future, Image-Scan/Paint is worth a look, as it integrates with Image-Read for OCR. You'll need Windows to run the Image-In suite.

If you don't require the general drawing and filling capabilities of paint programs, but only want a pixel editor to clean up scans prior to using them for DTP, then Optiks is good value for money. It has (very crude), pixel-editing capability, as well as other interesting image-manipulation features. More than this, it can import from dozens of graphics formats, and export to almost as many, enabling graphics format conversions. However, it is a budget priced Shareware program, so expect many bugs and limitations, (for example, the current version writes TIFF files without resolution information, making them unusable in programs like Designer). The full features of Optiks can only be obtained in the registered version — allow about two months after posting your order (and money), to receive the registered version from the author in the US.

In the field of OCR, Carets/Deluxe is the only TIFF compatible application tested. However, it is a program with only simple capabilities (as its price would suggest), in particular, being limited to a single font on a page with no graphics whatsoever. Within these restrictions, it does its job well. Again, if global TIFF compatibility is not an issue, Image-Read, at over three times the price, has at least three times the capability.

permanent swap file on the hard disk — it's a 300Mb unit, so a few megabytes for a swap file is neither here nor there. However, when I follow the instructions in the Windows manual for creating a permanent swap file, I get a message that the disk clusters are the wrong size, and that I can only use a temporary swap file, but the temporary swap file is a little too slow for my liking.

*Jan Harris
Moana, SA*

It sounds like you're using Wyse's own Dos 3.31, which is specifically designed for large disk partitions. This version of Dos (and similar ones from Compaq and others), got around the 32Mb limit by increasing the cluster size from the usual 2K, up to something like 20K, (in the case of a 300Mb partition). This saved the user's having to split the drive into 10 partitions, and worked with most software — Dos can support large partitions, however, its own format program can't normally create them.

In order to make the permanent swap file as fast as possible, Windows bypasses the standard Dos file operations, and works with the disk sectors directly. However, the problem with this approach is that the flexibility of Dos' cluster handling is lost, and to incorporate this flexibility into the driver for the swap file, would presumably slow down its operation.

There are really three solutions to this problem — or two solutions and a 'cop-out'. The cop-out is to put up with a sluggish temporary swap file. However, since Windows has to operate through Dos, and the file could possibly be fragmented, it is not very fast, as you have found out.

One possibility is to partition the hard disk so that you have one partition of the disk of 32Mb or less, so that the cluster size is the standard 2K, and use all or part of this partition for the Windows swap

Windows swap files

I'm running Windows 3.0 on my Wyse '386, and am trying to create a

SEPT 1990

file. The other is to re-format the disk under Dos 4.0, which lets you have large partitions and 2K clusters. However, this could cause problems with some disk utilities, so you might find that they need upgrading at the same time.

Hard disk maintenance

In Bruce Iliff's 'PC Maintenance' in June, he advises re-formatting a hard disk every six months 'to keep the disk in working condition'. If re-formatting is not carried out, how long can a hard disk be expected to continue reliable operation? If re-formatting is required, shouldn't a low-level re-format be required, rather than just a logical re-format? Digging for information on computers always uncovers more new questions than answers for old. Keep up the good work.

*Ross Byrne
Albany Creek, 21d*

It's true that it is a good idea to re-format your hard disk occasionally, although the reason given in the aforementioned article could be a bit misleading. The physical deterioration of the magnetic coating on the surface of the disk is not the real problem — it happens so slowly that other sources of wear tend to overtake its effects. However, there is a good case for periodically re-formatting a hard disk, to make sure that all the data on the tracks are aligned correctly.

When a hard disk warms up due to use, it starts to expand — the disk platters grow a little, and the head actuator arms get a bit longer. This causes the sectors written to a cold disk to be slightly out of line with those written when the drive is warm. That's why cooling is important — the hotter the drive gets, the more out-of-line the sectors become. Each track event-

ually starts to look like a higgledy-piggledy mess, and with normal wear and tear on the head actuator, the drive may begin to have trouble aligning the heads with some of the disk sectors.

Voice coil drives

I own an AT clone which includes a voice coil hard disk. I opted for this type of drive because of the extra speed over the stepper motor type. However, some people have told me that they are better than stepper motor drives because they automatically park the heads between read/write operations and when the power is turned off. Is this correct, and are there any other advantages of voice coil drives over steppers?

*Nick Trent
Wford, NSW*

The reason that most voice coil drives can automatically park their heads while stepper motor drives cannot, is related to the mechanical operation of the head actuators in the two types of drives. In the latter, the heads are moved across the disk surface one track at a time by a 'stepper motor', which is so designed that, when pulses are applied to its terminals in the correct sequence, the shaft of the motor rotates a pre-determined amount. So, for the heads in the drive to move inward four tracks, four pulses are applied to the motor; each pulse stepping the head inward one track. Because the amount through which the stepper motor moves when pulsed is known, the position of the heads can be determined by simply adding the steps applied to the motor in the inward direction, and subtracting the steps applied in the other direction.

A voice coil motor, on the other

hand, is a type of linear motor, not unlike the voice coil in a loudspeaker — hence the name. This type of drive is an analog system — it is theoretically possible to move the heads from track zero to the innermost track of the drive in a single action, rather than a large number of discrete steps. Thus, in a voice coil drive, it is a relatively simple matter to include a circuit which detects an oncoming power loss, and moves the heads to their landing zone before the power fails entirely. The power rails from the power supply take a finite time to collapse, thanks to the charge stored in the capacitors in the power supply — more than enough to move the heads of a disk drive to their landing zone.

However, such drives do not usually park the heads between read and write operations — to do so would seriously degrade the access time of the drive. The exception to this is the drives used in some laptops, which park the heads after a pre-determined period of inactivity, to protect them against the occasional jolt that is part and parcel of a laptop's lot. Some machines even power the drive down after 30 seconds or so, to save power, so of course, the heads are parked then.

You can get auto-parking utilities for stepper motor drives, so that if the power fails before you parked the heads, there is a good chance that they will already have been parked, provided you haven't used the hard disk recently. For a voice coil drive, there probably isn't any point, unless it is an odd one which isn't auto-parking.

Metric rulers for Corel

I've been using Corel Draw for technical illustrations — it's the best drawing package I've seen for a PC. While I think I've

mastered the package fairly thoroughly, I have never been able to figure out how to change the ruler from inches to centimetres. Because of the type of drawing I do, it hasn't bothered me (I just use the page size for reference), but a colleague has asked me. So, how do I change it?

*George Capon
Parkes, NSW*

Believe it or not, it's in the manual (at least it is in the version we have here — 1.21), not logically indexed, perhaps, but it's there. Under 'Show Rulers', it explains that the measurements in the ruler display are determined by the set of the grid size — change the grid to 10 divisions per centimetre, or whatever you want, and that's it.

My Dos manual only suggests the 'prompt' method.

*Richard Smith
Mundah, Qld*

Oops, sorry about that. Sure, the method to enter an escape character into a batch file is indeed simple, provided you know what it is! Assuming you are using Edlin, the escape character is entered by typing Ctrl-V, and then pressing the 'I' key — simple, if not a little obscure. Basically, the Ctrl-V character tells Edlin to replace the next character with the character that would be generated if that key were pressed in conjunction with the control key. So Ctrl-V A is Ctrl-A, Ctrl-V B is Ctrl-B, and so on. Ctrl-V I is Ctrl-I, which is identical to the escape key. This is actually documented in the Dos manual (in the section on Edlin), but thanks to its disorganised and poorly indexed nature, it's not that easy to find.

Other editors have different ways to put control characters into files. For example, the editor in the Norton Commander uses Ctrl-Q as a quote character, entering the next key pressed into the file literally. For example, 'Ctrl-Q Esc' will put the escape character at that point into the file. WordStar (in non-document mode, of course), is similar, except that Ctrl-P (yes, the printer control command), is the quote character.

So, to enter the ANSI code for flashing bright green text (the first example in the August column), you would type Ctrl-V[32;40;1;5m. The first 'I' is part of the escape character, while the second is the square bracket in the escape code.

October 1990

Desqview BAT files

I am using DESQview on a 20MHz '386 machine, and I have a query regarding batchfiles when invoked from it. When I terminate a COM or EXE file running in a DESQview window, the window closes automatically. However, when I quit an application started from a batchfile, I am returned to the Dos prompt. I then have to bring up the DESQview menu and select the 'Close Window' option. Is there any way to make these windows close in the same way as a COM or EXE window?

*J. Doyle
Glebe, NSW*

While DESQview can (and does), load COM and EXE files directly into memory and execute them, this does not extend to batch files, or indeed to COM and EXE files which do not have the full path and extension specified. In these cases, DESQview loads a copy of the command interpreter, command.com, to interpret the statements in the batch file. However, once a command interpreter is loaded, it will not terminate unless instructed, or it is killed off by DESQview.

To terminate a command processor, the 'exit' command is used. This can be typed at the Dos prompt after the batch file has terminated, or a better solution is to make the exit command the last line in the batch file. However, this still leaves you with a Dos prompt if the

batchfile is terminated prematurely (for example, by pressing Ctrl-Break), for some reason.

A far more elegant solution involves a completely different approach. Instead of entering the name of the batch file in the 'Program' field of DESQview's change program screen, put command.com, the command interpreter, there. In the 'parameters' field, put /c filename.bat, where filename.bat is the batch file to be executed. This method invokes command.com, and instructs it to execute the command specified on the command line after the /c switch, and then terminate, thus closing the DESQview window.

Filename.bat should be in the path, otherwise a full path name (including the drive) should be specified after the /c switch.

More on ANSI.SYS

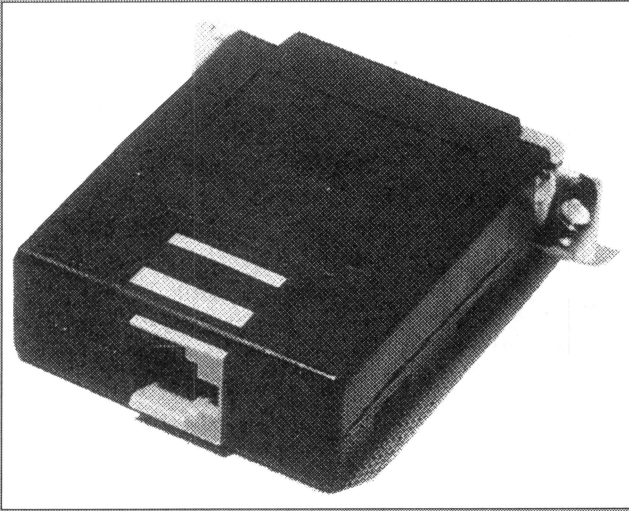
In 'TechTips', August 1990, you discussed ANSI.SYS, but I still have one question — how do you send the escape code to the screen so that Dos recognises it as the start of an escape sequence? I have tried typing ALT-27 on the number pad, Ctrl-I on the keyboard, IBH, and also using the escape key. Nothing works! From the way you say it, there must be a very simple way of passing this code on to the screen, but I can't find any way to do this in an echo statement, only as a 'prompt' statement.

Disk drive springs

I recently read in a book concerning maintenance of the IBM PC, that the gate

OCT 1990

Null modems



These RJ-11 adapters allow you to use modular phone cable for serial connections. They are supplied unassembled with push-fit pins attached to the modular jack. These are inserted into the appropriate holes in the DB-25 connector before assembling the complete unit.

ONE OF THE MOST USEFUL CABLES a computer communicator can have is a null modem. For those times when data needs to be transferred between two computers located relatively close to one another, modems are generally not required, and can be replaced by a hard-wired link, which in most cases can run much faster than a modem anyway.

Put simply, a null modem's job is to connect two computers together, in such a way that each computer 'sees' the other as a modem. The reason for this is simple — computer serial ports are designed to connect to modems, not other computers. If two computers' serial ports are connected together, nothing will happen, since the two machines will be transmitting on the same line in the cable.

A null modem swaps the transmit and receive lines between the two ends of the cable, so that the transmitted signal from one computer is sent to the receiving pin of the other computer's connector. If handshaking lines are involved, then these too are swapped over at one end. The use of handshaking lines in this way is a little non-standard, and different software will use the lines in different ways. However, in many cases, they may not be required, and a simple null-modem which just connects transmit and receive data, and ground, can be used.

If such a three wire null modem is all that is required, then a convenient solution is to use standard modular phone cable. Plug-in adapters are available from a number of sources, which convert standard DB-

25 connectors to RJ-11 modular phone jacks. These adapters are generally supplied unterminated, with pins attached to the six wires emanating from the phone jack. These pins are simply inserted into the required holes in the DB-25 connector. However, since modular phone cables typically have only four wires in them, the two outer ones can be cut off.

One point to bear in mind is that telephone cables swap the wires over from one end to the other. The best way to wire up such a cable is to use the two outermost wires as the two data lines, and both of the inside wires as the signal ground.

For the two data lines, insert the pins into holes 2 and 3 of the DB-25 connector. Push the pins in until they lock in place, and cannot be withdrawn. Make sure that the connections in each connector are identical, since the data lines will be swapped over in the telephone cable itself.

Now for the ground wires. Cut the pin from the end of one of the inner two wires (it doesn't matter which one), and bare the end of the wire for about 5mm. Then insert the bared end of the wire into the hole for pin 7 in the DB-25, and then insert the one remaining pin into the same hole, so that it holds the bare wire in place. This connects the two centre wires together, so that they both form the ground wire. If you are familiar with a soldering iron, then a neater solution would be to solder the bared wire to the top of the other pin, before inserting it into the hole.

If only one wire was used, the pin connection at each end would need to be different, which could cause problems if you end up making more than two of these connectors. It also means that one of the wires in the phone cable will be unused — which seems a bit wasteful, at the very least.

Once this has been done, the two connectors can be assembled, and you're ready to communicate. Plug a standard modular phone cable into each connector, and then plug them into the serial port on each PC. If you need extra length, then just buy a longer phone cable, or extend the existing one with an in-line joiner.

The DB-25 to RJ-11 adapters used come from David Reid Electronics in Sydney, (02 267 1385). They are also available from Black Box Corporation (03 725 2422, and Tandy electronics stores. Modular phone cables and accessories are available from virtually any electronics store, including those mentioned above.

on the disk drive should be turned down or shut when the machine is not in use. It was suggested that this would prolong the life of the springs in the disk drive, but I have not been able to confirm this. Can you shed any light on the matter?

Mark O'Brien
Bibra Lake, WA

That's a new one on me, although the explanation seems a bit backwards. On most floppy drives I've encountered, the springs are under tension when the door is closed, and relaxed when it is open, suggesting that spring life would be prolonged by leaving the door open! Also, closing the door places the two heads up against each other, which could be deleterious to their health.

On balance, I'd have to disagree with the suggestion in that book. In fact, when computer manufacturers ship new machines, they put a piece of card between the heads to stop them touching.

Memory query

I have been a reader of Your Computer for three years now, and a subscriber for almost as long. I enjoy reading the magazine for the variety of articles and the style in which they are written. Even though I am employed in the computer industry as an analyst, I still appreciate the effort taken by your writers to present subjects in an understandable manner, often starting at the basics.

I have a few questions which I hope you can answer for me. The first is about memory chips, specifically what size to choose, the type of chips (ZIP, DIP, or SIMM), and chip speed.

Secondly, my PC (a '286), has shadow video and BIOS ROM, and I am wondering what that means, and how it works. I have also encountered a slight quirk with Windows 3.0, when using it

with MS Word 5.0. I can start Word as an application from Windows, and work on a document. I can switch back to Windows in the normal way, by pressing Alt-Escape, suspending the Word session, and bringing the Windows screen back up. However, when I restore the Word session, by double-clicking on its icon, it asks me to 'Select text to copy formatting from', and I cannot move the cursor until I have pressed Escape. I am then returned to the command menu, and have to press Escape again before I can return to the text. This glitch always occurs after swapping the Word session out and in again, and happens on both my work and home machines.

J. Bruin
Glenbrook, NSW

We're glad you enjoy the magazine, and we intend to keep bringing you informative pieces about the PC industry and the technology driving that industry.

As a general rule, when looking for memory chips to expand an existing motherboard or expansion card, choose the largest size that the card can handle, as it usually works out cheaper on a per-byte basis.

However, make sure that any memory upgrades that you make are done in increments applicable to the bus width of the computer. So, an XT with its 8-bit bus has to be upgraded in 8-bit increments. In addition to the eight data bits, there is a parity bit, which works in the same way as a parity bit in serial communications — to detect errors in each byte.

So, for an AT or '386SX, the bus width is 16 bits, plus two parity chips, giving a total of 18 chips. A '386 has a 32-bit bus, so any memory upgrade will involve at least 36 chips. At least, it would if all memory chips were 1-bit wide. However, there are at least 3 types of memory chips in common use in PCs. The first of these are the standard 1-bit DIP animals, which

are still the most common variety on motherboards.

There are also 4-bit chips, which are common on video boards, which have no parity. However, when these chips are used on memory boards, they are used in pairs (to get 8 bits), plus a single 1-bit chip, for parity. The final type of memory packaging is a SIMM (single inline memory module), which is a tiny printed circuit wafer, with the memory chips surface-mounted on it. These then plug into special sockets on the motherboard, so that they stand out at right angles to it. There are two types of SIMMs around, 8-bit and 9-bit. The 8-bit ones are for Macs, which do not have parity-checked RAM, while the 9-bit ones include the parity bit which PC machines need.

Chip sizes are expressed in terms of the number of bits in the chip. So a single bit 256K chip is 256K bits, not 256K bytes. Nine of these chips will give you 256K bytes, with parity. A four bit 256K chip is 256K x 4, so two of them will give you a byte, and for parity, there needs to be a separate 1-bit chip. Of course, SIMM sizes are in bytes, since they are a byte wide.

The speed of memory chips is important too. If the chips are too fast, there is no problem, you'll just spend more money than you need to. However, if the chips are too slow, the data will not be ready when the processor expects it to be there, and errors will result. This is a major cause of parity errors in a computer. In general, you should use memory chips with the speed recommended for the system or board which they are to be plugged into. The speed required depends on many factors, such as the clock speed of the processor, the number of wait states, the type of processor, and whether or not a cache or other tricks are used.

Recently, one of the machines in our office, a 10MHz 0-wait state

AT, started to generate parity errors. Upon inspection of the motherboard, we found that the bottom 512K of memory was 100 nanosecond (about right for that speed machine), but the chips making up the extra 128K were only 150ns — a bit slow.

While RAM speeds have increased almost in step with processor speeds, the ROM (read-only memory), chips used for the BIOS on the motherboard, and things like disk controllers and video boards, are still relatively slow. So, any time a program needs to access one of these memories, the processor needs to wait for this slower memory, which slows the whole system down. Some systems allow these slower memories to be copied into spare space in system RAM, where they can be accessed faster.

This is known as shadow RAM, or shadow ROM, depending whether you are talking about the shadower or shadowee!

As for your Windows problem, we referred the matter to the technical support staff at Microsoft, who responded that it is a known bug in Word 5.0 when used with Windows 3.0. The bug has been

fixed in version 5.0b, which is available free of charge to registered users. Contact Microsoft for details on upgrading.

Boot-up menu

In 'Write Bytes', *Your Computer* June 1990, we suggested a simple way of having two autoexec.bat files, one which is executed when the machine is first booted for the day, and the other which runs on subsequent boot-ups. Of course, this can be extended to more than two, and you can have different config.sys files as well, with batch files to copy the appropriate one to the root directory before rebooting.

However, when you have more than a few different boot up options, this becomes messy, and if you need to change, say, your PATH variable, you have to do it in several files instead of just one. Also, when first turning the computer on for the day, you are restricted to the autoexec.bat and config.sys files which were left in the root directory when you last used the machine. If

you want a different one, you end up booting the machine twice.

Fortunately, there is a better way, thanks to a shareware program called Boot.sys, which allows a series of menu options to be placed in config.sys. Loaded as a device driver, Boot.sys lists your desired boot-up options on the screen, and then loads different device drivers depending upon your selection. For example, you can load QEMM for running DESQview, or himem.sys for windows. There are also options for a time-out and default selection, so that the machine will boot without operator intervention.

Boot.sys also passes the selected choice onto autoexec.bat, using the Dos errorlevel function, so different TSRs and programs can be loaded for each option. This system is better than having a menu in autoexec.bat, since the latter does not allow different device drivers to be loaded in config.sys.

Boot.sys is available from the Sydney PC Users' Group IBM board, on (02) 724 6813, as BOOTUP.LZH, and presumably other BBSes as well. And if you like it, don't forget to register!

Advice to give, or advice wanted?

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on. The best *Tech Tip* published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is Tech Tips, *Your Computer*, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

November 1990

BIOS drive tables

I recently purchased a 90Mb IDE hard drive from a well known supplier. The supplier implied that adding the drive to my computer would be a simple procedure, but as I have found, this is far from being the case. The problem stems from an incompatibility between the drive's configuration and the drive types supported by the BIOS in my PC (an Epson AX2). All of the drive types in the BIOS table are for 17 sector drives, while the new IDE drive has 33 sectors. I could set up the drive using a drive type in the table with only 17 sectors, but would end up with a 40Mb drive instead of the 90 I paid for.

The supplier said that the drive could use sector translation, where a drive type from the BIOS table with a similar capacity is selected, and the drive's in-built controller translates the BIOS calls to the real positions on the disk. While this works, the speed, as reported by CoreTest, was about 30 per cent slower than before.

I contacted Epson with my problem, and their technical support people told me that they could supply a customised BIOS with the correct formula in one of the unused locations in the BIOS table. This is the option which I ultimately pursued, although the new BIOS cost me \$90. It works extremely well. Is obtaining such a BIOS the only real way around this sort of problem, or is there another solution which I have missed?

*J Rowe
Arnccliffe, NSW*

The situation you describe is not uncommon, with drives starting to appear on the market with capacities far in excess of those which were around when the computer BIOSes were written. The BIOS drive table consists of a number of fixed-drive types, with pre-defined numbers of tracks, heads, and sectors. If the drive you're adding to the computer matches one of these built-in formulas, then there is no problem. However, if they do not match, a

number of possible solutions can be used to get around the problem.

The incompatibility arises not only with IDE drives, but also SCSI and ESDI drives. Some controllers have their own drive table on-board, which replaces the system BIOS table, and is specifically designed to support drives compatible with that controller. One ESDI controller which passed through our offices recently supported 256 different drive types!

However, if the controller (which is actually part of the drive in an IDE configuration), does not have its own drive table, other solutions need to be sought.

OnTrack's Disk Manager is one of the most popular software solutions, often shipped with non-standard drives. It allows large non-standard partitions to be defined on drives with track, head and sector counts different from those in the BIOS table.

A customised BIOS (such as you have obtained), is a more elegant solution, but you will need to go back to your dealer for this. Before going down this path, check to see if an updated BIOS is available which has the required drive parameters built in as one of the standard types.

A number of BIOS' also have a user-customisable drive type, which allows the relevant parameters to be entered from the setup program. Newer AMI and Phoenix BIOSes have this feature, but I am unsure whether Award BIOSes support it or not. With a user-defined drive type, the drive is selected from the setup program in the usual way, but then the individual parameters need to be entered in the table.

TIFF update

Further to Daniel Ford's analysis of

'TIFF-compatible' software in the September issue, we have subsequently tested Corel Draw!, version 1.21, with the same files, and found it to be compatible with all of them. So you can safely add this program to the list of true TIFF-compatibles. Oh yes, Corel Trace is also fully compatible.

Partitioning large drives

I am in the market for a drive of about 300Mb capacity, which I intend to use on my '386 under Dos 3.3. I do not want to upgrade to Dos 4.0 if at all possible, since it uses more memory than 3.3. However, a friend told me that you can only have four partitions on a physical hard disk, which means that I will only be able to use 128Mb of the space on the disk. I know some suppliers bundle a special version of Dos 3.3 with their machines, which allows large drives to appear as a single partition.

However, the Dos supplied with my machine is the bog-standard Microsoft product, which I understand cannot address partitions beyond 32Mb in size. Is there some way that I can address a 300Mb drive using standard Dos, or do I need one of those 'specials' that some suppliers have? I don't mind having 10 partitions on the disk, but I'm not too keen on spending money on a 300Mb drive if I can't even use half its capacity.

*Glen Myers
Gladesville, NSW*

Yes, Dos 3.3 can address a 300Mb hard disk, and use all of its capacity without waste. You appear to be a bit confused about the difference between disk partitions, and logical Dos drives. To address a hard disk of that size under Dos 3.3, you only need two partitions — a 32Mb primary Dos partition, and an extended Dos partition, using the remaining capacity of the drive.

NOV 1990

This is accomplished with the Dos Fdisk program.

Once the extended partition has been created, it then needs to be divided up into logical drives, again using Fdisk. The Dos 32Mb limit applies to the size of each of these logical drives, not to the size of the extended partition.

So, with that in mind, you can easily set up any sized hard disk under Dos 3.3, provided that the total number of partitions does not exceed 24, since Dos can only address 26 logical drives, and logical Drives A and B are devoted to the floppy drive(s). There are other limits to the maximum size of a physical hard disk, but they are imposed by the BIOS — the most commonly-encountered being the 1024 cylinder limit.

Losing at Reversi

A favourite pastime of hackers in days gone by was to play one computer off against another, or against itself if it was a multi-user system, at a game which was nominally a human versus machine game. By making the computer make the first move on one computer or terminal, and then entering the computer's move on each machine as if it was the player's own on the other one, the computers, unbeknownst to each other, are in fact playing another machine, and not a human opponent at all.

With Windows' multi-tasking capability, (even pre-3.0 releases), and good old Reversi, one can play the computer against itself at any desired level. However, in the window in which the computer plays first, the computer always loses. At novice level, it loses by 20 points, but even when the computer in the first window is

playing at expert level, the other window at novice level still beats it by two points! There is a limit to the times that a sentient being can do this without becoming completely bored, but it gives the brain a welcome rest after a hard day at the keyboard, (and mouse).

The downside of disk doubling

The following note comes from the Computer Products Division of Nashua, and concerns the practice of 'doubling' 3.5 inch double-density disks, to allow them to store twice the amount of data:

There is a theory around the computer industry that to convert a 3.5 DD diskette into an HD, one only needs to punch a hole on the other bottom side of the shell. Not only does this give a risk of leaving debris in the diskette, but a more significant risk is being taken by the user, in that the loss of data will occur.

Claims are being made that DD and HD media is the same and that higher prices are being charged for no apparent gain. We wish to point out to the industry that this is not so, and many users are being misled. The following information is offered as an explanation and should any user wish to query us further, we would be happy to discuss this in more detail.

Diskettes manufactured for high density applications are referred to as 3.5 inch DSHD (double sided, high density), MF 2HD (micro floppy, high density) or 1.44Mb formatted. Magnetic coating thickness is typically 1.7 microns and double density drives have 1.7 micron head gaps which is appropriate for the 2559 to 8717 BPI densities of MFM encoding.

The read/write heads of double density drives can create powerful magnetic fields which fully penetrate the 1.7 micron coating. However, the magnetic field generated by the narrow head gap of a HD drive may not fully penetrate the DD coating, leading to poor saturation and

associated performance problems. The field may also spread excessively in the thick coating, leading to excessive bit shift and data errors for high bit densities.

Initial results of testing were obtained using HD disks with an HD format in HD drive. As a result, amplitude performance was satisfactory. Resolution on track 79 was an acceptable 71 per cent at the Typical Operating Current of a stock drive and decreased to a marginal 60 per cent or less at higher write currents. Time Interval Analysis results were typical of good quality media and drives. Peak shift of the B6DB pattern is corrected by the drive controller during formatting, and the data bits of the formatted diskette are properly positioned in time. Window margins were 72 per cent or better, indicating that the disk should be readable with any good quality drive and controller.

When a DD disk was used with a HD format in the same HD drive, the results were very different. As a result, amplitude performance was unsatisfactory due to poor amplitude saturation. Resolution on Track 79 was an unacceptable 52 per cent at the Typical Operating Current. The time interval spectrum resembled that of MFM encoding but the peaks were excessively wide, leading to frequent read errors.

*Tony Little
Nashua Australia*

We've had our doubts about the reliability of formatting double density disks to high density ever since so-called disk doublers appeared on the market. It seems that because people have had some success in reading disks which have been formatted in this way, it is acceptable practice.

Certainly, for non-critical applications, such as transferring data from one computer to another, they would not cause much of a problem — if the receiving computer cannot read the disk, just try again with another disk.

But for important tasks, such as backing up a hard disk, using such

NOV 1990

Of TSRs and memory allocation

I RECENTLY OBSERVED THAT one of my colleagues had a fundamental misunderstanding of the way in which memory is allocated and used. While this is not important for people who only use one program at a time, it may make a great deal of difference to people using TSRs, and programs that dynamically acquire memory.

The case in discussion involves ISYS, a text retrieval program. The UPDATE module, which maintains the indexes, can make use of all the memory that is available, improving performance. My colleague also had the query portion, IQ, resident as a TSR. This takes approximately 110K of memory. With this, and Dos, and network software, he had just under 400K of memory. When the UPDATE starts, it senses the presence of IQ as a TSR, and asks if you wish to continue. The sensible approach is to stop at this point, return to Dos, unload IQ and then run the UPDATE. IQ can be unloaded from the ISYS main menu, or by calling up the program and unloading from there. ISYS even includes a command line parameter to unload the TSR. However, my colleague was remaining in UPDATE, calling up the TSR, and unloading it.

What he failed to understand, is that the memory freed was not allocated to the UPDATE program. The accompanying diagrams illustrate what happened: UPDATE cannot access the memory below it. The following three diagrams show the correct sequence of events. That is, unloading IQ at the Dos level, then running UPDATE. In this case, UPDATE is able to take advantage of the memory freed by unloading IQ. The extra 110K improved the speed with which the UPDATE completed, by about 10 per cent.

Inspired by this, I had a look at his environment. Despite the fact that he has a '386 with 2Mb of memory, he does not have an EMS driver installed. Most programs, ISYS included, do not make use of extended memory. We set up a controlled experiment, comparing UPDATE time on his machine, and mine, which is the same make, with 2Mb of memory and an EMS driver. On 'small' UPDATES, less than 10 minutes, the difference is negligible. However, 'large' UPDATES showed a 10 to 20 per cent improvement again. My colleague is a happy man. He is going to install an EMS driver, and has created a batch file that unloads IQ, and then runs UPDATE.

These are the relevant lines —

```

IQ UNLOAD
IF ERRORLEVEL 1 GOTO BADIQ
IDB UPDATE
IF ERRORLEVEL 1 GOTO BADUP
ECHO UPDATE completed successfully. ReLOADing IQ
IQ LOAD
  
```

GOTO END

:BADIQ

ECHO Unsuccessful in UNLOADING IQ.

GOTO END

:BADUP

ECHO Failure in UPDATE.

:END

The moral of this tale is that you should be careful in your management of TSRs. Those that can be unloaded, should be installed last, and they should be unloaded from the Dos command line, rather than within application programs.

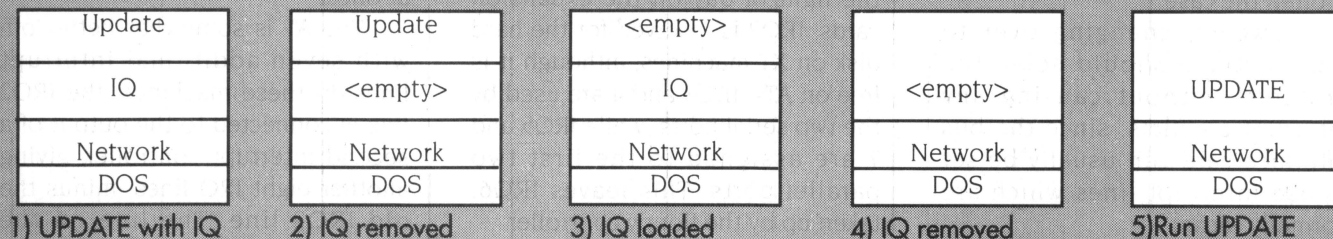
*Tim Keech
Ivanhoe, Vic.*

The above applies to any TSR program — if it is removed from memory, the memory which it occupied will not become available if other programs have been loaded into memory afterwards. This means any other TSR or application program, such as a Dos shell. Some TSRs will not allow you to unload them if something else has been loaded afterwards; others, however, will tell you they are unloaded, but there will be no change in the available memory.

SideKick is one of the worst-behaved programs ever created, and can prevent other programs being removed from memory even if the program in question has been loaded after SideKick. This is due to SideKick's grabbing the keyboard interrupt back from any other TSR which has subsequently taken control of it. When you try to remove the other TSR, the keyboard interrupt vector is pointing to SideKick, not to the other TSR, so it thinks that it wasn't the last thing to be loaded.

This behaviour causes no end of problems, and therefore, it is best to load SideKick after any other TSRs. If you need to remove a TSR, load it before SideKick, and then unload SideKick, then the other TSR, and then reload SideKick. More convenient is to use the public domain programs Mark and Release, which allow the removal of groups of TSRs automatically, using batch files. Try your local BBS for these.

When the IQ module of ISYS is unloaded, the memory freed is not allocated to the UPDATE program (this is typical of TSRs in general). These memory maps show what happened in the example described in the text —



disks is just asking for trouble. If the computer is lost or damaged (through theft or fire, for example),

the slightly different head alignment in the drive of a new computer could mean that the

already weak data patterns on the disk are totally unreadable — making the backup useless.

December 1990

Mouse choices

I currently have an AT with 2 serial ports, one of which is connected to a modem, and the other has a mouse attached. However, I have recently purchased an external fax modem for my laptop, and would like to use it with my desktop machine when I'm back at the office. The fax modem connects to the serial port, so I will need to add another port for the fax. What is my best option — stick with the serial mouse, and add an extra COM port (or two) for the fax modem, or put the fax modem on COM2, and replace the mouse with a bus mouse?

*David Nichols
Frenchs Forest, NSW*

While Dos and many application programs support up to four serial ports, there is a problem due to the lack of enough interrupt lines on the PC. To overcome this limitation, when four serial ports are installed, they share two of the interrupt lines, with the odd-numbered ports on IRQ4, and the even numbered ones on IRQ3. The problems involved were discussed in this column in the August 1990 issue.

If you were to add one or two extra serial ports to your machine, there is a good chance that an interrupt conflict would arise, especially if the fax software operates in the background, which is often the case.

However, changing over to a bus mouse should solve the problem, without causing any interrupt conflicts, since the bus interface card can usually be set to use interrupt lines which are normally vacant.

For example, the Microsoft 'InPort' interface card can be set to IRQ2 through IRQ5, inclusive. On most machines, one of these lines will be unused, and often there will be a choice of free lines. In your case, IRQ3 and IRQ4 are taken up by the two COM ports. However, IRQ5 is usually free (unless you have two printer ports), and IRQ2 is hardly ever used in an AT machine such as yours.

The only time you are likely to come across interrupt conflict problems with a bus mouse, is if you also have other 'non-standard' hardware installed, such as a LAN card or a scanner interface. In these cases, the only solution may be to use one of the serial port interrupt lines, in which case you will have to forgo one of the serial ports, leaving you no better off than a serial mouse.

This sort of problem doesn't just arise with serial ports and mice, but any device which needs an interrupt line. The trouble with interrupt-driven add-ons is that the number of interrupt lines in the PC is limited, and while the AT has almost twice as many interrupt lines as the XT, few cards seem to use them. The PC and XT only have eight interrupt lines, not counting the non-maskable interrupt, which is used to signal the processor that a parity error has occurred.

IRQ0 and IRQ1 are used for the timer and keyboard, and never see the light of day on the expansion cards. IRQ2 is reserved for the hard disk on XT machines, although it is free on ATs. IRQ3 and 4 are used by the two serial ports, while IRQ5 and 7 are assigned to the first two parallel ports. This leaves IRQ6, taken up by the floppy controller.

So, as you can see, on a full-populated XT, there are no spare interrupt lines. Most non-standard add-on cards therefore have a selectable interrupt level, so that if any of the interrupt lines are available, it can be used. For example, if there is no second parallel port, then IRQ5 is available for other uses.

The other solution is to share an interrupt line between two devices, which works with varying levels of success. On systems with four serial ports, they usually share the IRQ3 and 4 lines in pairs. This usually does not cause any problems as long as the software which is associated with each port is not run simultaneously. For example, if you have two different modems on COM1 and COM3, but only use one at a time, then there will probably be no problem.

However, if one of the ports is supporting a serial mouse, and the driver is resident when the other port is used, then a conflict will almost certainly arise. Most Dos software assumes that it has the interrupt line to itself, and gets horribly upset if this is not the case. There is no reason why devices cannot share interrupt lines, so long as the software's aware of this.

Another common trick is to allow a tape drive interface to share the same interrupt lines as the floppy controller, the assumption being that both are not used at once.

The AT is somewhat better off, with seven additional interrupt lines. In these machines, the IRQ2 line is connected to the output of a second interrupt controller, giving another eight IRQ lines, minus the old IRQ2 line. The line on the

DEC 1990

expansion bus which is IRQ2 on XT machines in reality is IRQ9 on ATs, but the BIOS re-directs this so that the software still sees the line as IRQ2. IRQ2 (or 9) is not used by the hard disk, leaving it free for odd-ball cards. IRQ8, 13 and 14 are the only reserved lines, used by the real time clock, maths co-processor, and hard disk, leaving four additional interrupt lines unassigned.

The difficulty with these extra lines is that they are part of the additional set of lines on the AT expansion bus, so that XT-compatible (8-bit cards) cannot access the lines at all. The only exception that I am aware of is the LANTastic interface card, which although it is an 8-bit card, has the extra 'finger' to connect with two of the interrupt lines in the extended bus, allowing AT machines to use one of these lines if available. It's a pity that more card manufacturers don't do this; it would make expansion of ATs much less troublesome, especially where lots of hardware has to be added to a single machine.

DeskJet draft mode

I have a Hewlett-Packard DeskJet printer, and when it is first switched on it defaults to letter-quality mode. However, the quality of this printer is so good that draft mode is adequate for most of my purposes. The letters are clear and a softer grey shade than the harsh black in letter quality mode. Also, draft mode is less demanding on the ink cartridge than letter quality mode. I can change to draft mode by pressing a button on the front panel, or by sending a code sequence to the printer from the computer. Is it possible to create a small COM program which can be called from Autoexec.bat to make the printer default to draft mode?

*J. Horn
Scone, NSW*

The simplest way to send a code to the printer would be to create a small file which is then copied to the printer, either when the system is booted up, or at a keyboard command. Unfortunately, Edlin cannot be used, because it will insert 3 characters after the desired ones — a carriage return/line feed combination, and an end-of-file marker, after the desired bytes. At the very least, it will mean that the first page printed after the code is sent will have an unwanted blank line at the start.

However, since there are only 5 bytes involved, debug can be used to create the file, using the following method. First enter debug in the usual way (check the Dos manual if you're unsure), then type in the following commands:

```
E 0100 1B 28 73 31 51
```

This instructs debug to enter five bytes into memory starting at relative address 0100, which is the location where debug will look when instructed to write out the file later on

```
Nfilename.ext
```

```
R CX
```

```
5
```

These two instructions tell debug the filename to write the bytes to, and the number of bytes to be written (five).

```
W
```

```
Q
```

These write the bytes to the disk file, and then quit debug.

So now you have a file containing the five bytes which need to be sent to the printer whenever it is desired to put it into draft mode. Now, put a line in autoexec.bat to copy it to the printer port whenever the system is booted up, or create a separate batch file expressly for the purpose. In either case, the statement:

```
COPY FILENAME.TXT LPTx
```

should suffice, where FILENAME.TXT is the file created using debug, and LPTx is the printer port to which the LaserJet is connected. If

you don't want the little 5-byte file cluttering up your root directory, it can be placed in a sub-directory, and the filename in the COPY statement changed to include the file's full path.

Fast format update

Thank you for an easy to read and not-too-technical magazine — it seems to be aimed at users like me who want to get the most efficient use we can from a PC.

I am having trouble with the program in the August 1990 issue, titled 'Fast format for Dos floppies', from J. Boetje. I believe I've followed the program to the letter, but it doesn't work. When I type FFORM A: or FFORM B:, the message 'File not found' appears. Have you any suggestions?

*R. Villarroja
Box Hill Sch, Vic.*

Oh dear, that one slipped through our testing. The 'File not found' messages are due to the script files, which should be telling debug what to do, not being in the current directory. Fast format will work fine, if it is invoked from the directory where the script files exist. The Dos path variable will let Dos find the batch file and debug itself, but is of no use for finding the two script files — FFORMA.SCR and FFORMB.SCR.

Fortunately, the solution is quite simple. Locate the two lines in the batch file which call debug and re-direct the appropriate script file into it —

```
debug < fforma.scr > nul
```

Then replace the name of the script with its full path name.

This should get the fast format program operating as it should. However, there is an important point to watch when using the program, especially in conjunction

DEC 1990

DEC 1990

with floppy disks that have bad sectors — in a word, don't! The fast format program writes over both the root directory and the file-allocation table at the start of the disk, which effectively renders the entire disk empty.

However, when the fast format program writes over the file allocation tables on the disk, in the process it destroys the bad cluster information contained in the FAT, so the disk appears to be 100 per cent error-free. That is, until you try to write some data to one of those clusters. So if you have any floppy disks with bad sectors, do not fast format them.

Second COM port

I have recently acquired a mouse, but when I connected it to the second serial port on my XT, it did not work. The software driver reported that the port didn't exist. The mouse works fine on COM1, but I have to unplug the modem to use it, which is a bit inconvenient. Upon closer inspection, I noticed that the I/O card has three empty IC sockets, one large one and two small ones. I presume that these are somehow related to the missing second COM port. With that background, my question is simple — what chips should I get to plug into these sockets, and where can I get them?

*Carla Guiterrez
Caulfield, Vic*

The three sockets you have noticed are indeed for the second serial port — many PC-board manufacturers have provision for the second port on-board, but don't put the chips in to save on cost. However, the chips are quite standard items, and easily obtainable from just about any computer store which sells add-on hardware.

The large 40-pin socket is for the ACE (asynchronous communications element) chip itself, the one which converts the parallel data from the computer to serial, and back again. The UART used in the PC and XT is a National NS8250, while that in the AT is the NS16450. As far as Dos and its applications are concerned, the two chips are identical in operation. You may as well just get the same chip as the other port on the card, for the sake of consistency.

The other two sockets accommodate the line driver and receiver chips, which translate the TTL levels used within the computer to and from the bi-polar RS-232 signals. Their numbers are DS1488 and DS1489, and can be obtained from many PC hardware suppliers, and most electronics component stores, (try local Dick Smith Electronics or Jaycar stores).

Once you have the chips, all you need to do is plug them into the board in the right places. If you haven't done this before, it would be a good idea to get somebody who knows what they're doing to plug the chips in.

That's all there is to it. Plug the serial board back in, and COM2 should be available. Depending on the particular card, you may have to change some jumpers — consult the card's manual for details on doing this, as they're all different.

Key codes for F11 and F12

In your discussion of ANSI.SYS in Tech Tips, August 1990, I note you do not give a function key code for the F11 and F12 keys. Somewhere I saw that their scan codes are 133 and 134, but although I have been successful in defining all the other function keys, there is no response from F11 or F12. Does this indicate that a new version of ANSI.SYS

is required, (mine is dated 17th March 1987)?

I am also having trouble with graphics under Basica. I have recently upgraded to a '386 with a VGA screen, and I thought a Mandelbrot set would look good on it. I tried using 'screen 4', but Basic reports this as illegal, although screens 0, 1, and 2 work fine. What am I doing wrong?

I also have an Epson EX800 colour dot-matrix printer, and have written Basic and Fortran programs which successfully output colour graphics to it. I also have a CAD package which can only produce hard copy output on a Hewlett-Packard plotter. Such a plotter is a little out of my price range at the moment. So I was delighted when I found the GRAPHICS COLORx command in my Dos manual, which should give me a small screen dump at the very least. Unfortunately, nothing happens! Is there any software package which will allow me to use my printer with this CAD package?

*E. Webber
Boroko, Papua New Guinea*

You're getting your money's worth, eh wot! Unfortunately, not many applications support the F11 and F12 keys, even though the extended keyboard layout has been around for several years. While it is easy to understand developers' reluctance to assign important functions to these keys, which don't exist on older-style keyboards, there's no reason they can't be used for user-defined functions, or shortcuts for more complicated key combinations.

Unfortunately, the BIOS call to access the extra keys on the extended keyboard, including F11 and F12, is different to the one normally used for keyboard access, which is why, for example, most programs don't see these keys and can't tell any difference between the left and right Control and Alt keys. To the best of my knowledge, none of

Perstor controller

I read in your 'Hard facts' article in the January '90 issue about the Perstor hard disk controller card, which can increase the capacity of a hard disk by 90 per cent. However, I've also seen a lot of horror stories about them on bulletin boards, and I'm wondering whether it is worth a try or not. I have a 70Mb Miniscribe voice-coil drive, and could use the extra capacity. Can you shed any light on this matter?

*D. Alexis
Chatswood, NSW*

Perstor controllers are one of those divisive issues, where people tend to be biased heavily one way or the other. Some users maintain that they are slow, unreliable and unpredictable, while those on the other side of the argument swear by them. I have been using a Perstor controller in an Everex AT, with a Miniscribe 3650 drive for 12 months now, with no permanent problems. I did have a minor hassle after installing the controller, with my AMI BIOS getting confused by the drive table which is on the controller card itself.

An updated BIOS, which incorporated a user-definable drive type, fixed the problem. The capacity of the drive jumped from about 41Mb to over 75Mb, which is now run as a single partition under Dos 4.01. The Perstor controllers achieve this feat by increasing the number of sectors per track from the usual 17, to 31, using a proprietary encoding scheme that the company calls ADRT — hence the 90 per cent increase in capacity. The number of tracks remains the same.

One of the oftmost-heard complaints about them is their lack of speed. Although the access time of the drive doesn't change, the data transfer rate does, as

there is now 90 per cent more data coming off the disk each revolution than before. So one would expect the data transfer rate to actually increase, and indeed it does. Sometimes.

The reason that many users (quite rightly), claim that the Perstor is slower, is that it doesn't incorporate a cache, as do some of the MFM controllers to which it is being compared. However, most of the MFM controllers which are being used do not have caching either, so such comparisons are unfair. Those using MFM drives really have two upgrade choices open to them: get more capacity with the Perstor, or a higher transfer rate with a caching controller. Perhaps Perstor will bring out a caching ADRT controller one day.

Of course, with ESDI, SCSI, and IDE drives, there is no competition, and these are still the drives of choice for high capacity and high-speed applications. However, these upgrade paths necessitate a whole new drive and controller, which is expensive.

The other criticism of Perstors is their compatibility with particular drive models. The controller comes with a list of compatible drives, but there are lots of (largely unsubstantiated), rumours around about even these recommended drives not working. However, I have not had any problems with my set-up, and Jerry Pournelle of Byte magazine swears by them.

I guess the safest way is to buy one and try it out, making sure that your money can be refunded if it is found to be unsuitable. If you can't get the Perstor to work, the worst you'll have to do is re-format the drive on the old controller, and restore all your data. If it does work (and on balance, I'd have to say there's a better than average chance that it will), then you've almost doubled your hard disk capacity.

the Ansi.sys replacements use this extended keyboard function either, so they too are oblivious to the extra keys. Until somebody writes an update of Ansi.sys, or one of its replacements, you won't be able to access the extra keys. If any of our readers do know of such a driver already in existence, please let us know, and we'll pass the info on.

The reason that your basic interpreter reports that screens 3 and 4 are illegal is that — the writers of GW-Basic, in their infinite wisdom, skipped over the values from 3 to 6 inclusive, so screen 7 is the next legal value, which has the same resolution as modes 1 and 2, but with a higher range of colours. For high-resolution colour graphics, screen 9 is the one to go for, while mode 10 offers high-res mono graphics. Screen 9 uses the full EGA resolution of the card, offering 2 or 4 bits of colour information, depending on the amount of memory on the EGA card.

As for using an Epson EX800 colour printer with your CAD package, there is a terrific utility from Insight Development Corporation, called PrintAPlot, designed for this specific application. It takes HP-GL files (a file created by plotting from the CAD package to disk), and rasterises them, converting them to a bit map suitable for printing to a wide range of dot-matrix printers, including HP Laserjet compatible models. It even includes a print capture utility, for packages which don't allow printing to disk. Yes, it does support the Epson EX800, and with colour printers such as yours, the pen information in the HP-GL file can be translated into ink colours, for full-colour reproduction.

Our copy of PrintAPlot came from Technical Imports Australia, (02) 925 0311.

DEC 1990

January 1991

Share and Dos 4.0

Occasionally, when booting up under Dos 4.01, I get a message warning me that 'Share should be loaded for large media'. My understanding is that share implements file and record locking for multi-user network use, and is unnecessary for stand-alone PC use. Can you please enlighten me as to the need (if any), for share on a non-networked PC?

*Nick Dyer
St Kilda, Vic.*

The share program in Dos 4.0 and later performs two functions. The first of these is the same as in previous versions of Dos — maintaining file and record locking in networked environments, as you already know. However there is another, less documented, function of share, which is important when running certain software on large partitions, (those greater than 32Mb in size).

When Dos was first released, the way in which programs interacted with disk files was through the use of file control blocks, or FCBs, which was similar to the CP/M method of file access. Later, with Dos version 2.xx and later, a new, more sophisticated method, was introduced, called file handles, which is the preferred method of accessing files on contemporary Dos systems. The number of files which can have file handles assigned (and therefore be open for access), simultaneously is determined by the FILES=XX statement in the config.sys file.

The trouble with the FCB method of file access, at least as far as Dos 4 is concerned, is that it knows nothing about large partitions, and so if you are running a partition with more than 16K of clusters, there's a very real possibility that the program will scramble some important areas of your hard disk in irretrievable fashion.

So in these instances, share performs a secondary role, translating FCB calls from software which uses them, to map correctly to clusters on the large partition. Because of the importance of loading share for large partitions, Dos takes some steps to load share itself if it can.

If there is a large partition present and share has not been loaded in either config.sys or autoexec.bat, then Dos will attempt to load it automatically. Dos uses the command interpreter to try to locate the share.exe file. So if your command interpreter (command.com) and share.exe are both in the Dos directory, then share will be loaded automatically without further intervention. But if you are in the habit of leaving command.com in the root directory, then share.exe must be located there also, leaving two extra files to clutter up your root directory.

There is no reason to have command.com in the root directory, and if you use the SHELL statement in config.sys correctly, then command.com can be loaded from the Dos directory, and provided share.exe is located there as well, it will be loaded automatically. The relevant line in config.sys should look something like:

```
SHELL=C:\DOS\COMMAND.COM
C:\DOS /P
```

The first argument instructs command.com as to where it should look when the need arises to re-load the transient portion of command.com. In Dos 4, it also uses that directory to look for share.exe.

If Dos can't find share, and you haven't loaded it manually, then the error message which you report will be displayed (rather fleetingly, sometimes) on the screen. However, even without share loaded, there may not be any cause to worry. Only those programs

which insist on using the old FCB way of accessing files will cause problems, and these are pretty rare nowadays.

This problem didn't arise with OEM versions of Dos 3.3, which implemented large partitions by using larger than normal clusters, as the actual number of clusters on the disk never exceeded 16K.

Dos 4 installation

I recently upgraded my computer from Dos 3.3 to Dos 4.01, using the automated installation procedure on the 'install' disk. However, one question that it asked me left me somewhat baffled, asking me to decide between program work space, and Dos functionality, or a compromise between the two. What do each of these three options mean, and which one would I select? I have selected the 'maximum Dos functionality' option, which seems to work all right for me, but I am not sure whether I've made the correct choice.

*Paul Stephenson
Leppington, NSW*

Basically, this part of the Dos install program is asking you what memory-resident software it should install for you in the config.sys and autoexec.bat files. However, no matter what option you choose during installation, you can change these two files at a later date to suit your particular requirements. Probably the best way to do this is to choose the 'maximum Dos functionality' selection, as you have done, and then edit these files to remove anything unwanted.

When you select the 'Maximum Dos functionality' option, you end up with config.sys something like:

```
BREAK=ON
COUNTRY=61,,C:\DOS\COUNTRY.SYS
BUFFERS=25,8
```

JAN 1991

```
FCBS=20,8
FILES=20
LASTDRIVE=E
SHELL=C:\DOS\COMMAND.COM
/P/E:256
DEVICE=C:\DOS\ANSI.SYS/X
DEVICE=C:\DOS\DISPLAY.SYS
CON=(EGA,437,1)
INSTALL=C:\DOS\FASTOPEN.EXE
C:=(150,150)
INSTALL=C:\DOS\NLSFUNC.EXE
C:\DOS\COUNTRY.SYS
```

and an autoexec.bat file along the lines of:

```
@ECHO OFF
SET COMPSPEC=C:\DOS\COMMAND.COM
VERIFY OFF
PATH C:\DOS
APPEND /E
APPEND C:\DOS\GRAPHICS
VER
MODE CON CP PREP=((850) C:\DOS\
EGA.CPI)
KEYB US,,C:\dos\keyboard.sys
CHCP 437
PRINT \D:LPT1
DOSSHELL
```

If you were to choose the 'Balance Dos functionality and program space' option, then you'd end up with only 20 buffers instead of 25, and no look-ahead buffers (rather than eight). Also, no FCBs (file control blocks) are installed, or rather, the Dos default of 4.0 is used. The first number indicates how many FCBs can be open at once, and the second number tells Dos how many FCBs it is not allowed to close itself. Dos does this if an application tries to open an FCB and there are no available FCBs, due to them all being in use.

FCBs are an old way of handling files. As described earlier under the 'share' heading, few programs now use them, and you can almost always omit this entry.

Also, the 'compromise' option will only install fastopen.exe for 50 files, instead of 150. Fastopen keeps track of the physical location

of files and directories on the hard disk, and saves Dos having to find frequently-accessed files each time it needs them. For each file that fastopen can track, it needs about 40 bytes of memory, so about 4K less memory is used by fastopen in this configuration than the one 'Maximum Dos' one.

All this saves about 13.5K of Ram over the first option.

The 'Minimum Dos functionality' option trims the autoexec.bat and config.sys files even further. No buffers are installed at all, and neither is fastopen. The maximum number of open files is eight. Also the graphics.com program (which allows graphics screen dumps to the printer) is not loaded, and neither is the resident portion of Dos' print spooler, print.com. This gives a further saving of memory of about 24.5K over the compromise option, or a total of 38K over the 'Maximum' option.

There is still a lot of useless stuff in there. For example, who ever uses Nlsfunc.exe? It is supposed to allow the use of country-specific information, but why anybody would need it is beyond me. Leave that out to save another 2K. Leaving out append.exe saves a further 10K — most programs are smart enough to find the files they need without append, which in my experience causes more problems than it solves. Display.sys is another driver that most people never need, saving another 18K of valuable Ram.

On the other hand, eight files is not really enough for most applications. Three files are always open — CON (the console), AUX (an alias for COM1), and PRN (the printer). The first one makes sense, but why on earth do PRN and AUX need to be open all the time.

Now, consider WordStar, for example. When you start the program (version 4.0), it automatically opens two of its overlay files — Wsprint.ovr and Wsmgs.ovr. As

soon as you open a file for editing, three more files are opened — the file itself, plus the two temporary editing files. That's eight files open in all, without trying anything fancy, like reading in another file. And WordStar is hardly a demanding program as far as the machine is concerned. So 20 files would usually be considered a minimum, and if you use a multi-tasking system such as Windows or DESQview, then there is a good chance you will need to be able to have even more files open at once.

Of course, if you have a '386 and 1Mb or more of memory, you can get the best of both worlds, by loading all of those memory-hungry drivers and TSR (terminate-and-stay-resident), programs between the 640K and 1 megabyte boundaries, still leaving as much space as possible for your applications.

And why the line 'Verify off' is added to autoexec.bat will probably go down as one of the not-so-great mysteries of personal computing. The 'off' state is the default state of verify, so there is no need for it at all. Personally, I don't mind the extra delay to make sure that what was written to the disk is what I thought would be written to the disk. Sure, you can use the '/v' switch after the Dos copy command, but try doing that while file-saving in your word processor!

Microbee to Mac

Firstly let me congratulate you on having consistently over many years produced a magazine that even someone like myself, who is at best, computer semi-literate, can enjoy and understand most of the material. I am hoping that you may be able to help me with a small problem.

After some six years, our faithful Microbees are now on their last legs, and, following your good advice, we have now

JAN 1991

JAN 1991

obtained a Mac Plus, principally because both my wife and I use Macs at work. Word processing is the primary use to which we put our computers, both at work and at home, and on the Mac we use Word 4.

My problem is that we have some 50 or so disks full of data from the Microbees, in CP/M format, using WordStar 3.1, which I would like to transfer to the Mac. Re-typing this volume of data is obviously out of the question. One suggestion that has been made is to have the data copied over to Mac-compatible disks commercially, however this appears to be a rather expensive operation.

In the October issue of *Your Computer*, you mentioned using a null-modem to connect two computers. Could you advise me if that is a feasible solution to my problem, and how I would need to go about solving it.

I know that on the system disk for the MicroBee there are a couple of communications programs, (it says so in the documentation). What sort of a program would I need for the Mac, and what sort of cables will I need?

Name and address supplied

A null-modem sounds like the solution to your problem, and you can either solder one up yourself, or buy one ready-made. A cable designed to connect a Mac to a serial printer with a DB-25 port (such as that found on the original ImageWriter), should work fine. Don't get a modem cable — it is not a null-modem, and won't work.

If you decide to make it yourself, we published details of a suitable cable in the January 1989 issue of *Your Computer* for connecting a Mac to a standard RS-232C port, such as that on the MicroBee. The wiring diagram is reproduced here for your convenience.

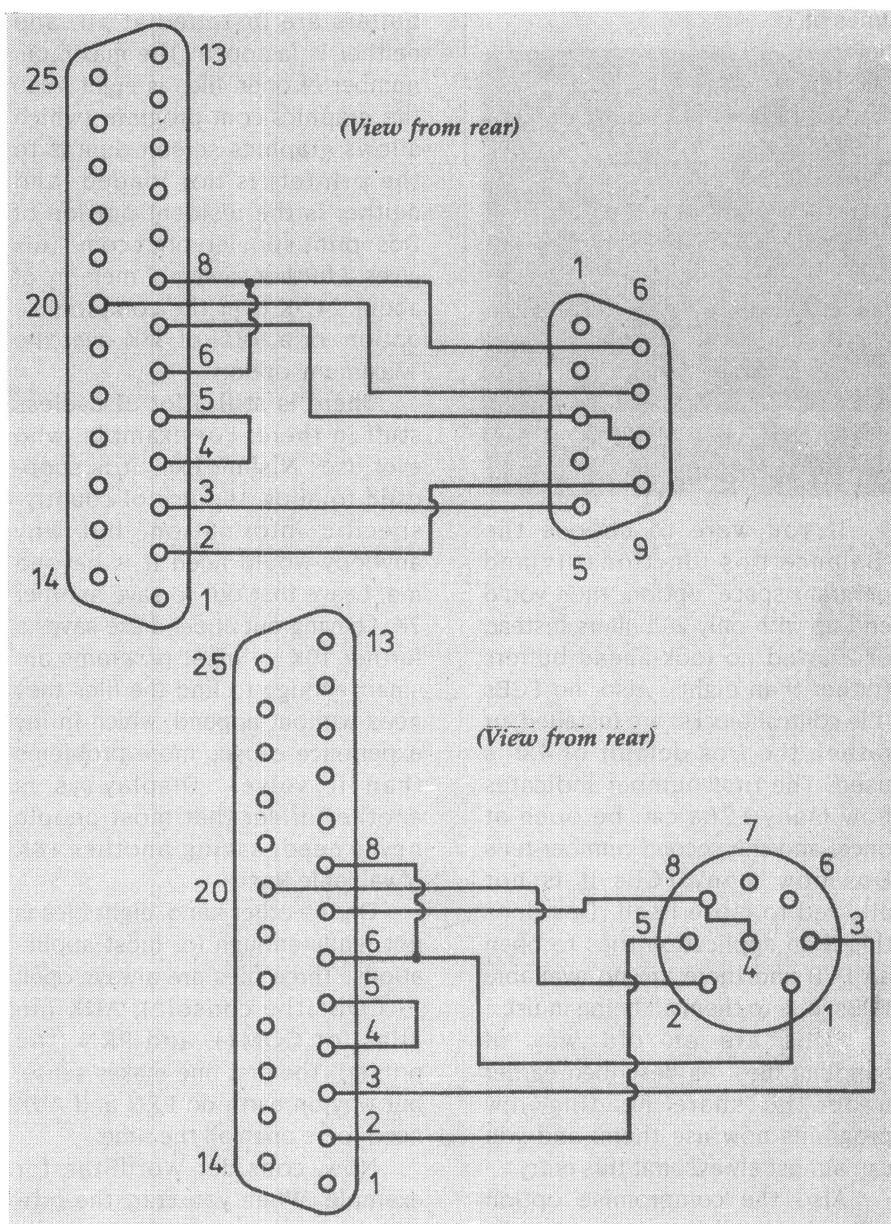
Once you have the cable, the question of suitable software arises. You need to have running on each computer communications software that both support the same protocol, such as Xmodem.

However, since you are just transferring text files (albeit in WordStar format) over a short cable, there should be no problem in using straight ASCII transfers, if that is all that the two comms programs have in common.

If my memory serves me correctly, the Telcom program supplied with the MicroBee supports XModem file transfers, and there are a few public-domain communications programs for the

Mac which also support that protocol — Red Ryder is probably the most popular. Your local Mac users' group should be able to help — the nearest one to you seems to be the ACT Apple User Group, PO Box 1231, Canberra, 2601.

Once you have software for both computers, it is now just a matter of setting up identical communications parameters at both ends. To retain WordStar formatting, you will need to use



This null modem cable (top) will connect a Mac to an RS232C equipped computer, such as the MicroBee; the same connection is shown (bottom) using the 8 pin mini-DIN connector.

eight data bits, otherwise the WordStar documents will end up as straight ASCII files on the Mac. The stop and parity bits can be set to 1 and 0, respectively. Now all you have to do is transfer the contents of the disks file by file to the Mac.

However, you will need some form of file conversion utility to convert the WordStar files to MS Word format. There are plenty of these around, but most of the ones which can read WordStar run on IBM compatibles, which is probably not a lot of use to you unless you have ready access to an IBM. Probably the easiest way around the problem is to convert the files on the MicroBee to straight ASCII, by opening them in non-document mode and typing Ctrl-Q U.

If you have access to an IBM PC, there is a more roundabout way that you can transfer the files, depending upon the software and hardware you have access to. FBN Software, (06) 285 2218, has a marvellous utility called PC-Alien, which can read and write a large number of CP/M disk formats, including MicroBee. We used this utility extensively ourselves back in the days when CP/M was in wider use than it is now. PC-Alien does not, however, perform any file format conversion — so the CP/M WordStar files will still be in WordStar format when transferred.

This would allow you to retrieve the files from the MicroBee disks and save them to Ms-Dos disks. The second piece of extra hardware you'll need is a DaynaFile drive for a Mac, which is capable of reading Dos disks, and copying the files on them to Mac format floppies or a hard disk. So while this is a two step process, if you have access to the necessary hardware it is a reasonably direct solution.

This still leaves the problem of file format conversion — WordStar 3.1 and Mac Word 4.0 are definitely problems, although they do share the ability to read and write plain

ASCII files. If you save the WordStar files in plain ASCII (non-document) format, and then copy the non-document files over to the PC, and then to the Mac, they can then be imported directly into Word. However, using this method, you will lose all type attributes, and there will be no distinction between soft and hard carriage returns, making the document appear as a lot of small (one-line), paragraphs, making justification and future edition difficult.

There are some utilities which can be used to convert word processing files; one of the best I've encountered recently is Word for Word, which is distributed by Mindscape, (02) 899 2277. This can translate between almost any current word processor, including WordStar and MS Word.

Once the files have been converted to Word format on the PC, then they can be transferred to the Mac via the DaynaFile drive, and read into Word.

Volume labels

Why does Dos 4 ask me for a volume label whenever I format a floppy disk, irrespective of whether I have specified the /v switch or not. More to the point, how do I stop this annoying behaviour?

As most people would realise, in Dos versions before 4.0, format did not ask for a volume label when formatting floppies, unless specifically asked to, through the /v switch. Under Dos 4.0, the function of the /v switch has changed slightly. Rather than specifying whether or not to put a volume label on a disk, it is used to specify a volume label on the command line.

So format a: /v:label formats a disk in the A: drive, with a volume label of 'label'. Unfortunately,

leaving a blank after the '/v:' switch results in an error message. One solution is to either label all disks with a dummy name, such as 'no-label', or to use a blank character, such as ASCII 255 (entered by holding down the Alt key, while typing '255' on the numeric keypad) after the /v switch.

The result of this is a volume label which prints as a blank, rather than reporting 'No Label'.

All this is pretty keystroke-intensive, so you might want to create a batch file for formatting disks in this way. By the way, this Alt-255 character can also be used in filenames, if you want to keep their contents away from prying eyes. Give the file a name less than eight characters long, and then use Alt-255 as the last character of the filename. Since that character prints as a space, it will show as a blank, but whenever anybody tries to edit the file, they won't be able to, unless they know this trick too!

Watch out for programs like the Norton Commander, which displays the files on the screen, and allows them to be edited by simply pointing to them with a highlight bar. Since the Norton Commander 'knows' that the funny character is there, it will have no trouble opening the file to edit it.

Parity errors

Sometimes when I boot up my computer I get a message saying something like 'Parity Error', whereupon the machine hangs, and needs re-booting. When I first noticed the problem, re-booting was all that was needed to fix it, but lately it seems to have been getting worse, and I have to re-boot several times before I can use the computer at all. Also, it occasionally happens while I am using an application, which is inconvenient, to say the least.

JAN 1991

The computer in question is a President 10MHz AT, running at zero wait states, with 1Mb of Ram. I originally bought it with 640K of Ram, but a friend upgraded it a little while ago to 1 meg, using some chips he said he had 'lying around gathering dust'. Thinking this could be the cause of the problem, I had a look inside the box where the Ram chips are (according to the manual), and have located two sets of 18 identical chips, although I can't tell which are the new ones and which are the old ones. One set of chips (in two rows) has the marking 'TMS41256-15NL' on them, while the other chips are similar, but the number ends in '-10'.

Is it the new Ram chips, or something more sinister, and just a coincidence that the problem has occurred now? Should I get my friend to put the old Ram chips back in, and leave me with 640K. I'd rather stick with a full meg if possible, since I have just bought Windows 3.0 (after reading about it in your July 1990 issue), and it is a bit crippled with only 640K of Ram.

*Phil Ashley
Moana, SA*

From what you write in your letter, I'd say your diagnosis is correct. The number after the dash on those Ram chips indicates their access time, in tens of nanoseconds (ns), so the chips ending in -15 are 150ns chips, while the others are 100ns. For a '286 processor running at 10MHz, 150ns is a bit slow for zero wait state operation — even 120ns chips are a bit borderline, although I have seen similar machines to yours running 120ns chips, with no trouble.

The parity error arises from a feature built into all IBM PCs, from the very first one — Ram parity checking. Instead of using eight chips for each bank of Ram, IBM (and of course, the cloners) used nine — eight to give the eight bits of storage in each byte, and a ninth, which is used to keep track of the accuracy of the data stored in the

Ram, by storing a parity bit for each and every byte in the memory. When a byte is read, the parity is checked, and if there is a discrepancy, an interrupt is generated. The trouble with parity checking is that although it alerts you to a problem, it can't tell you which bit is in error — it could be any of the eight data bits, or the parity bit itself — so the only thing to do is alert you to the problem, and halt the system.

In your case, the parity errors are arising because the processor is trying to read the data from the Ram chips before they are ready. When this happens, the data from the Ram chips is not valid, and any discrepancy between the actual data on any one bit, and the data that was originally stored in the chip, will cause a parity error.

So you really have two choices — put the original Ram chips back, and lose the extra 384K of extended memory, or buy some faster 41256 chips — 100ns ones. Now that Ram prices have dropped back to more sensible levels, this would not be a major expenditure.

Windows icons

You may or may not know that you can have files represented as icons in program manager groups under Windows 3.0. To achieve this, maximise the Program Manager, and then fire up File Manager over the top of it. Re-arrange the windows so that you can see the desired Program Manager group underneath the File Manager Window. Then locate the file which you want to 'iconise', and simply drag it into the Program Manager group. It's that simple.

However, this will only work if the extension of the file in question appears under the (Extensions) heading in the Win.ini file. The icon

of the associated application program is then used as the icon for the data file. If the extension of the file does not appear in Win.ini, then you will not be able to drag it into the Program Manager window.

Say, for the sake of argument, that you want to put your autoexec.bat file in a program group, so that when you double-click on it, Notepad fires up and loads autoexec.bat for editing. First, copy the autoexec.bat file to autoexec.txt (the actual name doesn't matter, nor do the contents of the file, but this is convenient), and then drag the new file to the Program Manager.

Once that has been completed, the autoexec.txt file can be deleted. Now select the icon for the new file (under Program Manager), and then select 'Properties' from the file menu. Change the 'Command line' entry to read:

NOTEPAD.EXE C:\AUTOEXEC.BAT

and then click on OK. Next time you double-click on this icon, you'll be able to edit the file instantly.

Internal or external modem

I want to buy a 2400bps modem for my laptop (a Toshiba T1200), but I'm not sure whether to get an internal or an external model. An internal one would be nice, but they seem to cost a fair bit more than external ones. I am also hoping to get one with inbuilt fax capabilities, if this is possible.

*Jan Matheson
Sutherland, NSW*

There is no definitive answer to this one — it depends. Most laptops have an internal slot for a modem, but these are usually proprietary, so in a lot of cases

DOS 5.0

AS IF DOS 4 wasn't enough trouble, Dos 5.0 is definitely on the way, having already completed its Beta testing, although Microsoft is being typically vague about release dates.

Probably the most useful feature of Dos 5 is its ability to load itself into high memory — the first 64Kb of extended memory. Of course, this will only work on AT and better machines, and of course, only those with extended memory available.

For those who still use Basic, the QuickBasic interpreter will be bundled along with GW-Basic. This will allow users to write better-structured programs than Gee-Whizz allows, and leaves the way open to compile them (by buying the full QuickBasic package) later on if they accidentally turn into programming masterpieces.

DosShell has also received a facelift, resembling Windows' File Manager, so that OS/2, Windows and Dos will all have a similar user interface. Of course, Dos applications will still perform their own screen management, so this resemblance will only be apparent when Dos isn't doing anything.

Limited on-line help for all Dos commands is also included, and can be accessed by typing `?/` after the command on the command line. An optional command-line editor is also included (about time!), expanding on the rather crude facilities provided by the function keys under previous versions.

Support for 2.88Mb (ED) 3.5-inch floppies is also provided, which should coincide nicely with the availability of suitable drives from Toshiba, and probably others. The disks will no doubt cost a bomb for a while though.

Dos 5.0 might be available by the time you read this, but if not, it can't be too far off.

you are restricted to the modem manufactured by the same company as the laptop. With popular models such as the Toshiba which you have, and the fact that Toshiba has the same modem slot in most of its machines (except the new notebooks), there's a number of suitable third-party modems.

Since you are also looking for fax capability, which the Toshiba modem does not have, one suitable model for you would be the WorldPort 2496, distributed by Dataplex. It supports modem communications using both Bell and CCITT frequencies at up to 2400bps, and Group III fax at a full 9600bps. Since you have a Toshiba, you have the choice of either an internal or external unit. TouchBase Systems, who make these modems, also have a model for the new notebooks, but that one is only available in the US at present. Dataplex can be contacted on (03) 735 3333.

Internal modems have the obvious advantage of not contributing to the space occupied by computer equipment when travelling — fitting as they do within the case of the laptop itself. However, if you use an internal modem in a battery-powered machine, such as your Toshiba, then it will contribute to the drain on the battery, although most laptops can power down the modem when not in use.

Also, if you use a mouse with your laptop, an internal modem will leave the serial port free for this, so you don't find yourself constantly swapping cables and installing and un-installing mouse drivers.

The other problem arises when you want to upgrade to a new machine at a later date. There's no way you're going to get a Toshiba-slot modem into a Compaq, and even the new Toshiba notebooks have a differ-

ent slot (although again, it is the same design for several models). And don't expect an internal modem to add much to the re-sale value of the machine. If a potential purchaser doesn't want a modem, there's no way he or she will pay extra for one.

External modems for portables generally cost less than internal ones, which is the opposite to the case with desktop PCs; the reason is that the proprietary nature of internal laptop modems naturally limits them to smaller production runs than for standard ISA-bus modems, so initial development costs have to be recouped from fewer sales.

Also Austel, Australia's telecommunications regulatory body, has rather more stringent (some would say inflexible), requirements than similar bodies overseas, so the cost of getting a modem up to Austel specifications and actually approved becomes significant when you are only expecting to sell a few hundred units.

But if you want to use it with more than one computer (there is nothing stopping you using it on the desktop at home as well), or if you think you may be getting a new laptop in the future, with a different modem slot, then an external model is the one to go for. They generally run on internal nine volt batteries (the ones we called 'transistor batteries' in the transistor radio days), but even the alkaline ones only last for a few hours at best, due to the high current drain of the circuitry.

There is a 2400bps modem on sale in the US, which draws its power from the laptop and phone line, but I think the chances of it gaining Austel approval are pretty slim. The extra size of a portable modem is not really a major problem when travelling, as they really are quite small now — certainly smaller than any desktop model.

JAN 1991

VGA on old monitor

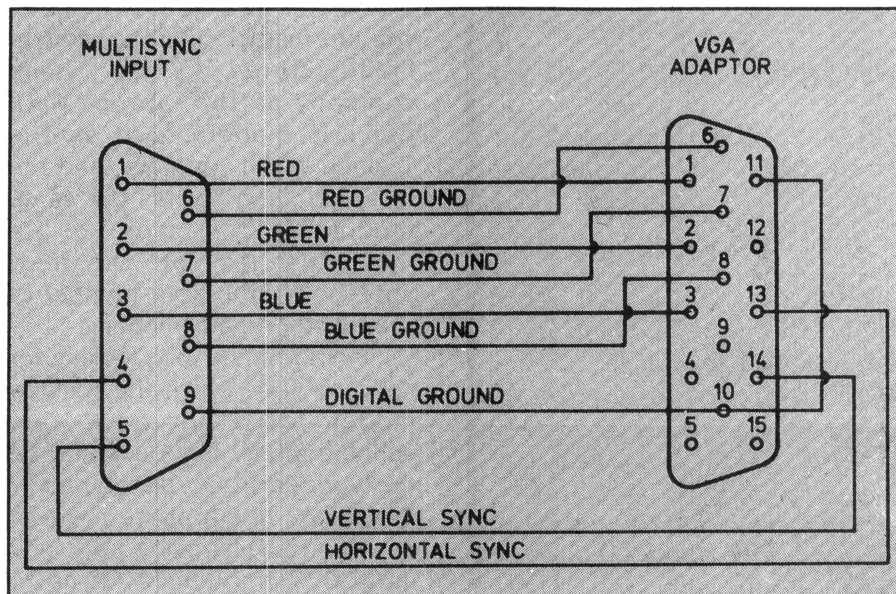
I'm in love! I've just seen Windows 3.0 running on a VGA, and I want one. My current system has an EGA card with an EIZO 8060S monitor. According to my manual, the monitor can display analog signals (which, I gather, is the reason VGA can display so many colours), but I am unsure whether VGA signals are special in some way — the manual only mentions PGA, which I've never heard of. Can I get a VGA card to work with this monitor, or do I need a new monitor as well?

*Greg Bouman
Paisley Park, Vic.*

We presented an article in our April 1989 issue, describing how to connect a pre-VGA multi-sync type monitor to a VGA card. To be able to do this, the monitor needs to be capable of handling analog inputs, rather than the digital TTL signals used by the EGA and previous adapters. From your description your monitor sounds like a suitable candidate.

The only thing you will need is a suitable adapter cable. This cable needs a 9 pin D-connector on one end, to match the one on the back of the monitor, and a miniature 15 pin connector on the other end, to plug into the connector on the VGA card. Don't use the 9 pin connector found on a lot of VGA cards — this is a TTL output, and although it will work, you won't be able to see the full range of colours that the VGA standard is capable of. Ready-made adapters are also available from various computer retailers, if you aren't the soldering iron type.

The design of the particular monitor will determine how well it copes with the VGA signal frequencies. For example, the NEC Multi-Sync (the original model), does not perform vertical auto-sizing, so you will have to adjust



The pin connections for Multisync input to a VGA adaptor — note the jumper between pins 10 and 11 on the VGA adaptor. If a shielded cable is used, connect the shield to the 'digital ground' pin at each end of the connector.

the vertical height control when changing modes, or put up with a slightly squashed-looking text mode.

On the other hand, some later models work perfectly well without this manual intervention, and your Eizo could be one of those. Also, some monitors will handle the SuperVGA mode supported by most contemporary cards, with a resolution of 800 x 600 pixels, so you can go even better than VGA. Some cards can display even higher resolutions (up to 1024 x 768), but the chances of displaying legible images at this resolution are pretty slim. You really need a 16-inch or bigger monitor to do this.

programs into memory above the 640K limit. However, try as I might, I cannot get it to work. The supplied 'Optimize' program crashes whenever I run it, and if I try to use `loadhi.com` or `loadhi.sys` myself, they report that there is no high memory. By the way, my machine has 2Mb of Ram, and I am assuming that some of that is high Ram and the rest is ordinary extended memory.

*K. Foster
Jindalee, WA*

Self-installing software has some disadvantages, one of which is that many software publishers use it as an excuse for poor coverage in the manual — 'Just put the disk in the A: drive and type 'install'.' Unfortunately, many machines seem to have compatibility problems with the installation program, even though the software itself may run fine. However, manually installing some software is impossible, while with others, you often need to see a successful installation on another computer, in order to see how to install it yourself. Not much help if you only have access to one computer. QEMM's

QEMM crashes

I recently purchased the Quarterdeck Expanded Memory Manager, with the main intention of using it to put device drivers and terminate-and-stay-resident

Optimize program is one such program, although the problem is related not to Optimize, but to QEMM itself. The manual method is explained in the manual, but is spread across several sections of the text, and the relationship between these sections is not obvious at first.

Before looking at the specifics of your problem, it is useful to understand just what high Ram is. As you already know, conventional Ram ends at the 640Kb boundary, but the addressing range of the 8088, and the '386 in real mode, extends to 1Mb. In this space between 640K and 1M, the PC designers located various other memory devices, primarily the BIOS Rom, and Video RAM. However, there are also quite a few gaps in this space, which is usually unused, except by network cards, and a few other add-ons. Unfortunately, this unused space is not usable on 8088 machines, since there is no memory located in these gaps. However, the memory management facilities of the '386 allow this chip to make some of its extended memory fill these gaps. Of course, this reduces the pool of extended memory available, but proper use of high Ram can increase the amount of conventional memory available to applications.

DESQview, however, does not automatically set aside memory for high Ram — you need to specify this on the QEMM command line, using the RAM switch. This is the cause of the 'no high RAM' message which you have seen. The QEMM Optimize program puts this switch on the command line, and then re-boots the system. When you specify 'RAM' on the command line, QEMM looks to see what regions of the memory map are available, and maps extended memory into these locations, making it into high Ram.

Unfortunately, the method used by QEMM to work out which regions of high memory can be used as high Ram is not fool-proof, and it will occasionally put memory in locations which it should not. QEMM cannot tell what memory locations are likely to be used by software which hasn't initialised yet,

and will sometimes put Ram in there, causing a clash (and crash), later on.

The solution is to use the optional parameters for the Ram switch, which tells QEMM which areas of high memory it can put Ram in, and which areas to leave free. To do this, you need to know what areas of memory can be used for high Ram, and QEMM comes with a utility (called Qemm.com) which tells you which areas of memory can be allocated as high Ram, and those which cannot.

Firstly, install QEMM in your config.sys file, without the RAM parameter, and then re-boot the system. If you have a network, make sure that its drivers are also loaded, so that memory locations used by the network card do not get allocated to high Ram. Then type 'QEMM T' at the command line, and it will show up any unused areas of the memory map as 'mappable'. This display will be more reliable than the automatic method used by the QEMM device driver when it is loaded, since the operating system is fully initialised, and the Qemm.Com program can more accurately determine which areas of memory are free.

Make a note of those memory regions above 640K (A000) which Qemm.com says are either mappable or rammable — for the purposes of high Ram, they mean the same thing. Each of these regions should then be specified on the QEMM command line in config.sys. For example, if C800 to DFFF is mappable, then put the switch 'RAM=C800-DFFF' on the command line. Put one of these 'RAM=...' switches for each block of mappable Ram which you want to use as high Ram. Then re-boot, and re-run QEMM.COM, and those regions of memory will then be reported as high Ram.

You can then use Loadhi to put device drivers and TSRs into this high Ram, as described in the manual. By the way, I like to use the 'best fit' (/B) switch, which causes loadhi to put a driver or program in the smallest block of Ram in which it will fit. This means that all the small things will fill up the

Hint

WORRIED about PC security? Putting a password program in your autoexec.bat file is fine, but easily circumvented by booting from a floppy disk in the A: drive. However, if you only have one floppy drive on your controller, there is a simple fix.

Simply unplug the ribbon cable from the back of the drive, and plug it into the connector that would normally connect to the B: drive. The drive is now the B: drive, and any attempt to access the A: drive will result in an error. You will probably need to set some DIP switches (on an XT) or run the setup program (on an AT), to recognise the drive in its new location.

Of course, you can now no longer boot from a floppy disk, so make sure you are careful with changes to the config.sys and autoexec.bat files. Using this technique, and locking the keyboard (if the BIOS allows the system to boot with a locked keyboard) allows an unattended computer to be protected from all but the most determined villain.

small blocks of high Ram, leaving the larger blocks free for bigger programs which need the extra space. Loadhi cannot split a program across memory blocks, it can only load them into an unbroken block of high Ram.

Also, bear in mind that many programs use more Ram when initialising than they do in their resident state. So if you have a 64Kb block, that doesn't necessarily mean that two 32Kb TSRs will fit in it. There's nothing you can do about this — if there is not enough high Ram for a program to initialise, then you will just have to put it in conventional memory. However, if you do have problems fitting everything into high Ram, put the smallest possible things into conventional memory, so as not to waste too much of it.

JAN 1991

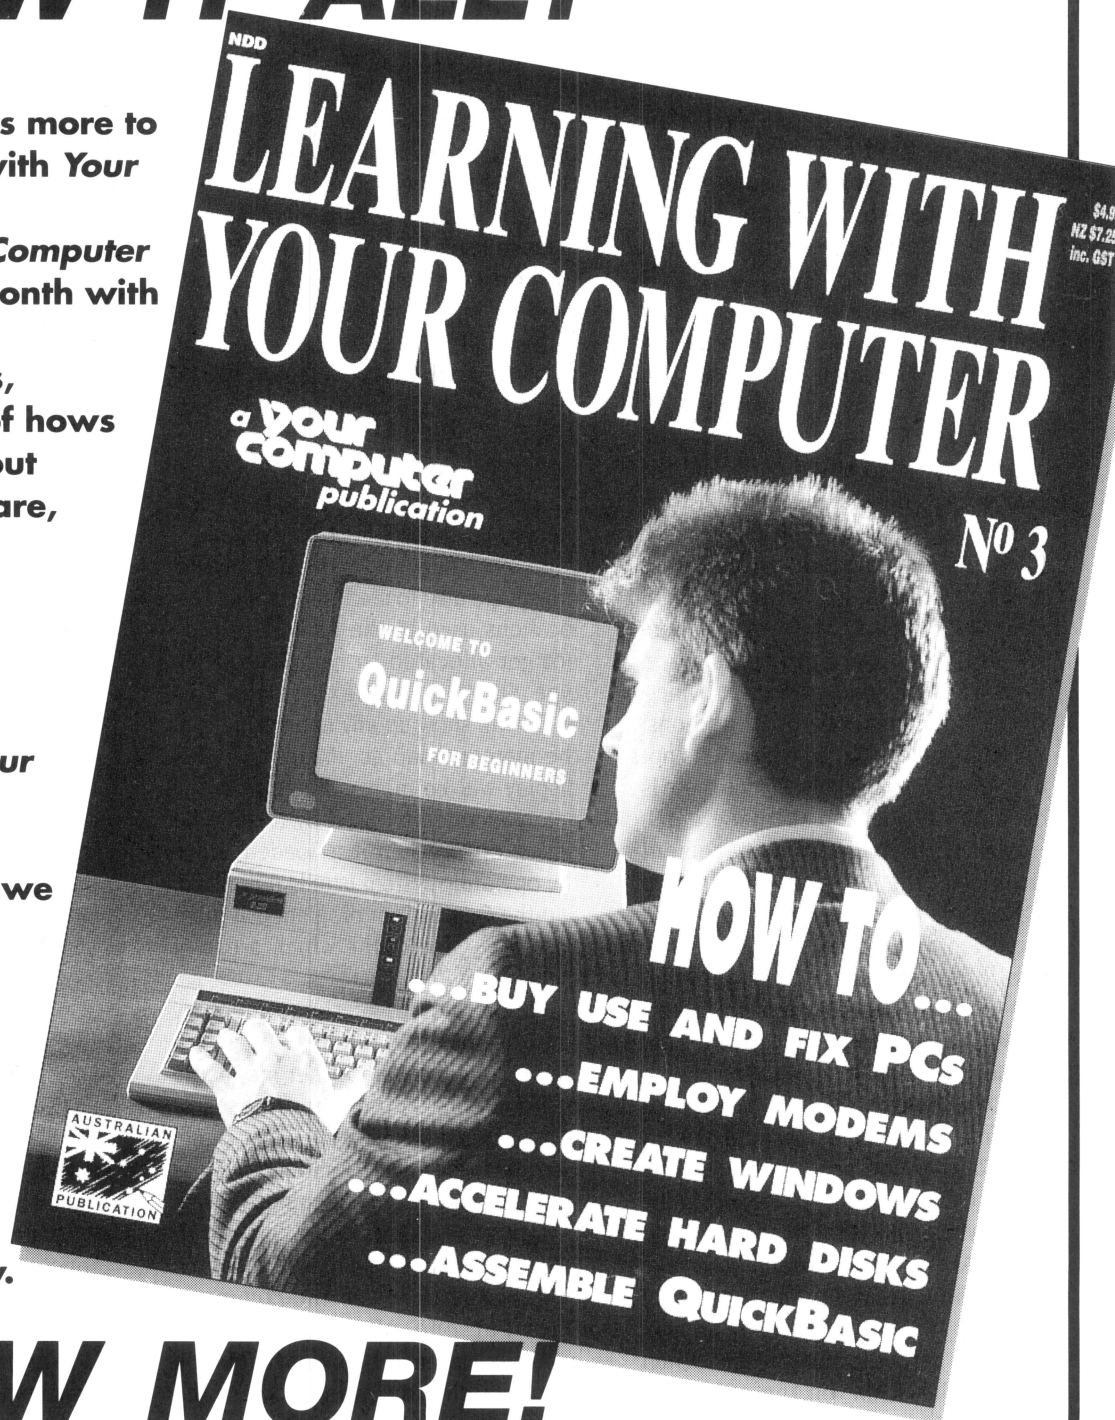
KNOW IT ALL?

There's always more to be learned with *Your Computer*.

That's why *Your Computer* is packed each month with news, views, opinions, reviews, tutorials, — lots of hows and whys, — about software, hardware, peripherals, applications, multimedia, communications, programming. In short, *making your micro work*.

And periodically, we publish *Learning with Your Computer*.

Learning with Your Computer Nos 3 and 4 are on sale now at your newsagency.



KNOW MORE!

Learning with Your Computer No 4 has features on multimedia, on- and off-line mass storage, CD-ROM, modems, and a hugely comprehensive optical storage media primer. There's plenty on C++, Forth, and the concluding tutorials on Assembling QuickBasic, (begun in *Learning with Your Computer* No 3).

IMPORTANT NOTICE: *Learning with Your Computer* Nos 1 and 2 are completely sold out.

ON SALE NOW

February 1991

Motherboard memory

I have a rather elderly AT (about 3 years old), which I put together myself from boards acquired from a variety of sources. The motherboard's RAM uses 41256 chips, giving me 1Mb of storage, but I would like to increase that to 4Mb, without using an extra memory board, if possible. Can this be done? Presumably I will need 411000 chips (if that's the right number), but the motherboard manual only mentions 4164 and 41256 chips.

E. Catling

The simple answer is 'No', it can't be done. Not easily, at least. The motherboard was obviously designed before 1Mb chips became widely available, and so no provision was made for them. So, unless you feel like cutting tracks and soldering wires to the motherboard, forget upgrading it, and get yourself a memory board.

There is another advantage of using a memory board, especially on a '286-based system. Most boards can be configured as either extended or expanded memory, so that you can select the 'appropriate type of memory for the applications which you are running. While EMS emulators are available, which make extended memory on a '286 look like expanded, they are slow, since the processor needs to be running in protected mode to access the extended memory, and switching back to real mode necessitates resetting the processor, which wastes a considerable amount of time.

Interlaced video

I have a Trident TVGA SuperVGA card, and an NEC MultiSync 3D monitor,

which I use to run AutoCAD at 800 x 600 dots resolution — the extra resolution is really useful for examining detailed drawings. However, upon closer inspection of the documentation for both the card and monitor, I discovered that they both support an even higher resolution mode — 1024 x 768 dots.

Unfortunately, whenever I switch to this mode, the screen flickers noticeably, to the point of being virtually unwatchable, especially when displaying shaded regions. Is there some mutual incompatibility between the high resolution modes of the card and monitor, or is there another reason. At the moment, I am limited to 800 x 600 mode, which is why I originally bought the card and monitor, but if it is capable of displaying a higher resolution, why not use it?

J. Brian

The flicker which you have observed is caused by the way in which the card and monitor generate the image for display, called interlacing. An interlaced display consists of two fields — one containing the odd-numbered scan lines, and the other containing the even-numbered ones. The complete screen image consists of both of these fields, which are displayed alternately on the screen, rather than the entire image being displayed in one sweep of the electron gun.

Monitors have a specification called bandwidth, which describes how fast the electron gun can be turned on and off, and consequently how many pixels can be displayed in one sweep of the electron gun. The higher the bandwidth, the higher the resolution of the image, all other things being equal.

Interlacing is a means of doubling the resolution of the display, without needing to double the bandwidth of the monitor, which is expensive. Since, on each sweep of the screen, the electron gun only builds up half the picture, the amount of info that can be displayed on the screen is effectively doubled.

We all know what they say about free lunches, and monitors are no exception. The price of this doubling of screen information is that any given pixel is only updated half as often as it would be on a non-interlaced display, so it appears to flicker. High-persistence phosphors on the back of the screen can minimise this problem, but they are generally only used on monitors designed for interlaced display, rather than universal monitors such as the MultiSync. The IBM 8514/A monitor uses interlaced display, and IBM claims that there is no discernible screen flicker.

Ordinary TV pictures are also interlaced, but the flicker is not noticeable, due to the much higher screen to eye distance.

Aside from getting a high-persistence multisync-type monitor (if such a beast exists), the only real solution is to get a non-interlaced monitor (such as the NEC 4D), and a card capable of driving it. The other consideration here is that most 14-inch screens do not have a fine enough dot pitch to display 1024 x 768 pixel images — interlaced or not. The MultiSync 4D has a 16-inch tube, giving that much extra display area for complex images.

Hard disk connector

I recently came across a Priam 803-11 110Mb SCSI hard drive. It was originally housed in a mini computer made by a company called Tecmar. Unfortunately it appears to have a non-standard power connector, having six pins in two parallel rows of three. After contacting the Australian distributor for Priam, I was informed that they had gone bust, and that the distributor had never heard of the 803-11. I would be most appreciative of any information that you could come up with.

I have only just begun reading your magazine again after a two-year break. I

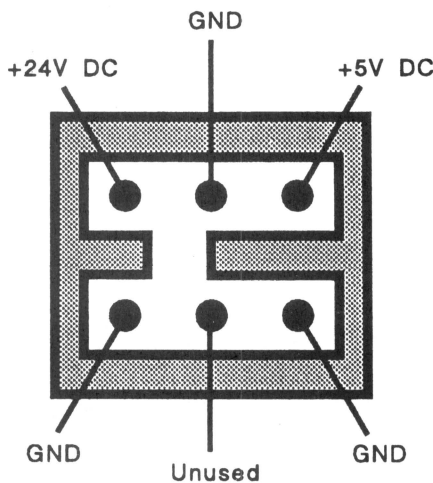
FEB 1991

FEB 1991

was a most avid reader back in the days when 32K of Ram in a Z80 or 6809 was a big thing. You appear to have changed things a bit, but you're still looking good.

*Chris Dann
Vermont, Vic.*

Thanks for that, and welcome back! While I am not familiar with the particular model of drive either, it sounds to me like the type of power connector used on 8-inch disk drives in the days when they were a little more common than they are now. These drives had a +24V motor supply, and a +5V logic supply, and were powered through a 6-pin connector similar to the one you describe. This is how the connector is configured —



Just for safety, take a look at the drive motor to see if you can ascertain its voltage rating, and make sure that it really is 24 volts, or you could do some expensive damage. However, I think it is unlikely that Priam would have used a 12 volt motor on a drive with a 24 volt power connector. If you are unsure, you could always try powering it from 12 volts at first, and if the motor doesn't spin up, then try 24V.

Backing up several PCs

In our office we have four PCs, which we use for various tasks. The machines are not networked, since we seldom have any need to share or exchange data between them, since they are all used for different purposes, and a simple printer switch box allows us to share printers.

However, following a disk crash on one of the machines, it was decided that we should invest in some form of backup system, probably a tape drive. However, adding a tape drive to each machine seems a little expensive, and installing a network seems a bit 'over the top' just for backup purposes. Ideally, what we want is a system similar to the printer sharing box, where a single tape drive can be shared between the machines, and switched between them as required.

R. Joyce

Tape drives for PCs can be either internally or externally mounted. Internal units are cheaper, since they save by not needing a case or power supply.

However, a single external unit with several controller cards will usually cost less than an internal drive for each machine. A controller card is installed in each of the computers, and when it is desired to back up the drive in a given machine, the tape drive is plugged into that machine, so the backup can be performed.

While plugging and unplugging the tape drive cable might be a bit inconvenient, it is a lot less trouble than feeding a stack of floppies into the drive one at a time, or re-keying a few megs of lost data. External tape drives generally use 37-pin D-connectors, and obtaining a 4-way switch box to switch this many lines could prove difficult. Also, long cables may not be very reliable at the sort of speeds used by tape drives, so a switch box would be of little use unless the

machines were located close together.

Don't forget, too, that if you were to install a simple network (such as LANtastic), it would not only allow you to back up all hard disks using a single internal tape drive, but would also eliminate the need for your printer sharing switch, since the network would take care of allocating the printer to users as they needed it.

Parallel to serial

Our office recently upgraded from a rather ancient CompuPro multi-user CP/M system to a Novell network of PCs, and this included upgrading our printers to an HP LaserJet series III, connected to the server. This works fine, but we have been left with two orphaned serial printers, a Toshiba dot-matrix, and a Brother HR-15 daisy wheel. Both of these printers are still in good condition, and would be useful to take some of the load off the LaserJet, and to serve as a backup, should the laser break down for some reason.

The server has two serial ports and two parallel ports, with one serial port connected to a modem for remote logins to the network, and one of the parallel ports connected to the LaserJet. This leaves us with two serial printers, but only one serial port, and the dealer who sold us the network says that Novell cannot support a third serial port. Is there any way that we can connect one of these printers to the spare parallel port, or will we just have to switch the spare serial port between the two of them?

G. Carr

Black Box has a bewildering range of 'this-to-that' boxes, which can convert just about any interface to any other one, (well, almost). The company has a converter called the Serial — Parallel Converter II, which should do exactly what you

want — allow you to use a serial printer on a parallel port (it also does the reverse). Just buy one of these and connect it between the parallel port and the printer.

They can also be supplied with a 256K or 512K buffer, but since you are presumably using Novell's built-in spooler, there would be no advantage in your case. Black Box can be contacted on (03) 725 2422.

Printer driver problem

I have been a bit hesitant in writing to you, as I am not a regular subscriber; I buy my magazines from the local newsagent and as such, do not buy every issue of Your Computer. Nevertheless, I am hoping you will be willing to help me through your 'Tech Tips' column.

I have a copy of MS Works 2.0, which I use for word processing, and an HP DeskJet Plus printer. I have purchased 2 proportional-spaced font cartridges for the printer, and am having trouble with the output which Microsoft support has been unable to resolve (although they have been helpful). They have provided me with a number of drivers, called DeskJet 1, 2, 3, and 4, that are also used for MS Word.

The problem seems to occur (generally, but not always) if the 14th character position from the left margin is a space. Drivers 1 and 3 both add two extra spaces at this point (giving three in total), while drivers 2 and 4 omit the space and join the words together.

I am presently using driver 4, and having to proof-read each letter, then add an additional space where words are joined together, before printing the final copy. The real frustration begins when I then edit the letter, or change margins or indents in any way, as I then have to proof read again, add additional spaces, and remove those spaces previously added.

P. Cain

Unfortunately, having neither an HP DeskJet printer nor MS Works readily available to me, it is impossible to try to simulate your problem. If you have access to another printer (not an HP DeskJet), you could narrow down the problem to the driver, by printing a problem document to this printer, through the appropriate driver. I assume that you have no problems with the printer with other applications — it would be wise to check, to be sure.

Assuming the above confirms that the driver is at fault, another thing you could try is using the drivers for the HP LaserJet, as the DeskJet is supposed to emulate that printer (although I'm not sure which model). If that doesn't work,

perhaps you should ask for your money back.

Laptop battery capacity

I have a Toshiba T1200 laptop with a hard disk, but I cannot get it to work for any length of time from the battery. Normally this is not a problem, as I use the machine in the office and keep it plugged into the mains, so that the battery stays charged up. However, I took the machine out the other day, and it only worked for a few minutes before the battery failed. Fortunately I took along the power supply, and was able to continue working from the mains. The machine is a couple of

How many files?

You've just bought yourself a new laptop, and you're trying to install your favourite software on its hard disk. Out with Flight Simulator 3, on its two 360K disks. The laptop only has a 3.5-inch drive, but fortunately your main machine has both flavours of drive, so you proceed to copy the two distribution disks onto a blank 3.5-inch diskette. The first disk copies without drama, as does most of the second one. However, before the copy can complete, the message 'File Creation Error' appears. That's odd; the contents of two 360K disks should fit onto a 720K number without any trouble, a directory of the target disk shows there is still space to spare.

The cause of this problem is quite simple, really. When Dos formats a disk, it creates (amongst other things) a root directory on the disk, which stores filenames, attributes, file length, and the starting cluster of the files and directories on the disk. However, this root directory is fixed in size, and on a double density floppy

(360- or 720K), it can contain only 112 files or directories. Since there are more than that many files on both Flight Simulator disks, there isn't room in the root directory of the 3.5-inch disk for all the files from both source disks, even though there is enough space for the files themselves. High density (1.2- and 1.44Mb), disks can have double that number of files in their root directories, and hard disks can have 512.

The solution? Create a sub-directory, and put all the files in there. Sub-directories can expand as needed — in fact they are just special files as far as disk allocation is concerned, and Dos simply adds another cluster to the end of the directory to put more files in when required. So, provided there is enough space on the disk itself, you can put as many files as you like into a sub-directory.

By the way, don't forget that a system disk has two hidden files on it which add to the total, and if you give the disk a volume label, that counts as another file.

FEB 1991

years old, but I would have expected a better life expectancy from the battery than this.

S. O'Toole

These are the classic symptoms of the 'memory' effect, which plagues nickel-cadmium batteries, such as those in your Toshiba. When the battery is repeatedly discharged, but not totally, it tends to remember the amount of charge delivered, and on subsequent cycles the battery will only deliver this amount of charge to the laptop before going flat.

At first glance, it might appear that you shouldn't have a problem, since the battery hasn't been discharged at all, with the power for the machine coming exclusively from the mains adapter. In truth, the battery has been cycled every time that you have used the machine. Nicad batteries, when

left alone, tend to self-discharge, at a rather high rate — this can be as high as one per cent per day. All batteries do this to some extent (even car batteries), although most are not as bad as NiCads.

So each time the machine is turned off, the battery starts to discharge, and this discharge will be enhanced by the Toshiba's resume mode (if you have enabled it), which maintains power to the system memory when the power is off. Some later model Toshibas do this whether resume mode is enabled or not, to power the 'hard Ram' virtual disk drive, although I am not sure about the T1200.

When the machine is powered up again, (say, the next day), the charger starts charging the battery again, but because the battery wasn't fully discharged (in fact, it was nearly fully charged, the memory effect starts to become noticeable, especially if this cycle occurs regularly, like every day. So when you tried to use the laptop without the power adapter, the battery could only supply as much power as it had during those small overnight discharges, and so it died after a few minutes.

Fortunately, the problem is not permanent, and can be reversed (at least partially) by charging and discharging the battery for a few complete cycles. I'm not sure whether this will ever fully restore the capacity of the battery, but it will certainly come close. Plug the charger into the laptop, and let it charge the battery fully up, (leave the machine turned off if possible, as this will hasten the charging process). Then unplug the charger, and run the laptop from the battery until the battery goes flat, and the machine shuts itself down. Repeat this cycle a couple more times, and each time you should notice an improvement in the time that the laptop functions before the battery fails.

For the curious, I have never come across an adequate explanation of why the memory effect occurs, at a chemical level, and that's not for want of trying. However, it seems to effect small NiCads more than larger ones, and very large nicads, such as those used in solar power systems, do not seem to suffer from memory effect at all. Not much comfort for those of us who can't carry around a 20kg battery with our laptops!

Phone line sharer

Like many users, you probably only have one phone line to share between the phone and modem. However, it's more than just a tad annoying to have somebody pick up the phone in the middle of a file transfer, which usually results in the transfer being aborted.

Arista Electronics has a handy little device, which allows two devices to share a phone line, but when one is using the line, the other is cut out. Although it is intended to allow two phones to connect to a single phone line, it can in fact be used with anything which connects to a telephone line, including modems, faxes, and answering machines. So just connect this little gadget to the phone line, and plug the phone(s) into one socket, and the modem into the other.

Now, whenever you are using the modem, any phones connected to the other socket on the sharer will be dead, and if a phone is in use, the modem cannot grab the line and disrupt the conversation. Incoming calls are diverted to both sockets, and the first device to answer (phone or modem), gets sole use of the line for the duration of the call.

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on.

The best *Tech Tip* published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is *Tech Tips, Your Computer*, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

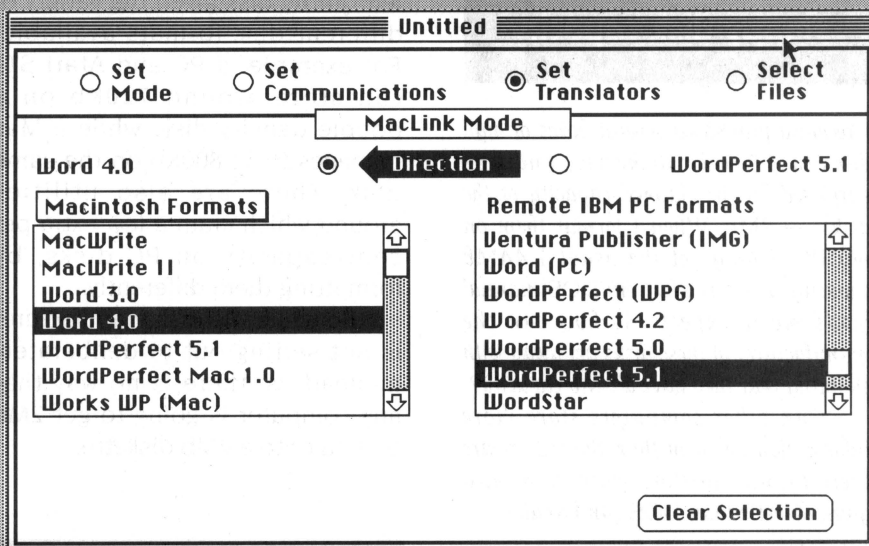
From PC to Mac and back

MANY READERS will be familiar with Traveling Software's LapLink — LapLink III being the most recent, and most versatile, release to date. For those who haven't seen or used it, LapLink is a rather unusual communications program that allows files and directories to be copied from one PC to another, via the machines' serial or parallel ports. LapLink III is one of the most-used packages in the *Your Computer* offices, enabling us to install our favourite applications, utilities, and benchmarks on review machines in a single operation, without the tedious swapping of floppy disks back and forth.

MacLink Plus/PC follows a similar idea, but across platforms, specifically between the Mac and PC environments. As any Mac user knows, the physical formats of Mac and MS-Dos floppies are as different as chalk and cheese, and until Apple started putting FDHD 'super drives' in their machines, there was no way that you could take a disk from a standard PC drive, and put it into the Mac, and read it. The fact that most PCs had 5.25-inch drives didn't help much either — folding the disk over on one edge so that it would fit into the Mac's drive was seldom successful, and tended to make re-use on the Dos machine difficult as well. Transferring and translating files between platforms account for a lot of the enquiries we receive for the 'Tech Tips' column, and of course with the IBM PC and Mac platforms being the two dominant PC platforms, this package should solve a large proportion of those problems. As a bonus, it can also transfer files between the Mac and either a Sun or NeXT workstation.

Sure, there were third-party solutions, such as a PC expansion card, which connected to a standard Mac external floppy drive, or the DaynaFile, which connected to the Mac's SCSI port, allowing 3.5- and 5.25-inch Dos disks to be read and written as 'normal' Mac storage devices. But they were expensive, and didn't address the question of file-format translation. Some packages such as MS Word and Aldus PageMaker use the same file formats on both platforms, but this is something of an exception. Also, while an application might be the program of choice on the PC platform, its Mac version might fall short in terms of capabilities of other equivalent Mac applications.

Enter MacLink Plus/PC, from DataViz Inc. In a similar way to LapLink, it is a



hardware/software solution, consisting of a set of disks, and a few cables. Obviously, since MacLink is a cross-platform file transfer facility, software is supplied for both the PC and Mac. The PC software is also supplied on both 5.25- and 3.5-inch diskettes, while of course the Mac only needs 3.5-inches.

The main cable is a PC-to-Mac null-modem, with a DB-25 connector on the PC end, and an 8-pin mini-DIN plug for the Mac end, which can be plugged into either the printer or modem ports on this machine, while the PC end is restricted to the first two serial ports. A 9-pin adapter ensures that PCs with these smaller serial ports are covered as well. An odd omission is any provision for connection to earlier Macs, which used DB-9 connectors for its serial ports. The third component of the cable assembly is another (though shorter) DB-25 to 8-pin adapter, which allows MacLink to link two Macs, or link a Mac to a NeXT workstation.

MacLink Plus operates at any standard serial port speed from 300- (ouch!) up to the default 57,600bps, which is as fast as the Mac can go. Firing up the software at both ends and giving the 'connect' command gets things going, with the transfers and translations being controlled from the Mac end.

The first thing to do is to select what file translation (if any), is to be used. For example, a WordStar file from the PC can be converted to MacWrite, or MS Word on the Mac. The available Mac file formats are displayed on the left-hand side of the Mac screen, while those for the PC appear on the right. Once the file formats have been

selected, and the desired direction for the translation is chosen the files may be selected.

When the 'select files' option is chosen, the same two panels which displayed the available file formats now display the files on the selected drives on each computer. The Mac's hierarchical file structure can be navigated in the usual way, while the PC's directory tree also appears in the form of a Mac HFS.

Multiple files can be copied in a batch, and MacLink automatically chooses filenames appropriate to the destination platform, although these can be changed if desired.

MacLink Plus/PC supports literally dozens of file formats, including word processors, spreadsheets, databases and graphics applications. A total of 31 Mac formats and 37 PC formats can be converted, in addition to 10 formats each for NeXT and Sun machines. While there are some conversion limitations when special features are unsupported in one or the other of the programs, MacLink does its best to translate as much information as possible.

MacLink Plus/PC also has a 'local desktop mode' which just performs the file translation function, with both the source and destination files already residing on the Mac file system. This allows MacLink's powerful translators to be used if you have a DaynaFile or FDHD drive, or a mixed PC/Mac network, thus bypassing the time-consuming serial file transfer process.

MacLink Plus/PC is distributed in Australia by Infomagic (02) 975 1044, and is priced at \$330.

FEB 1991

Floppy disk capacity

I recently purchased several boxes of high density 3.5-inch diskettes, and was surprised by the claimed capacity of the disks — 2Mb. When I format them on my PC, I only get the usual 1.44Mb capacity, which is what my DOS manual leads me to expect. So how can the manufacturer of these disks get away with claiming that they have a 2Mb capacity?

Can other computers store more information on them than the PC, or are there utilities to store more data on a given diskette than DOS can handle?

G. Townsend

No, the manufacturer is not taking you for a ride! The capacity of the disk is indeed 2Mb, unformatted. However, the process of formatting the disk uses up some of it. The sectors need to be marked out on the tracks, and other space is taken up with Cyclic Redundancy Check (CRC) information, which ensures the integrity of the data on the disk, (and is used by some of the more insidious copy-protection schemes). The file allocation table and root directory also take up some of the available space.

Once all of this formatting information is taken into account, the capacity drops from the claimed 2Mb, to 1.44Mb. As far as the disk is concerned, it is storing 2Mb of data. The fact that some of the data is simply there so that the computer can find the other data is irrelevant as far as the disk itself is concerned. Equally irrelevant to the end-user is the 2Mb figure.

So unless you can store your information on an unformatted disk (and if you can, we would like to know how you do it), the unformatted figure is not a lot of use. One reason manufacturers do display the unformatted capacity is that they have no real way of knowing how much space will be

left on the disk given the variety of different disk formats available. For example, a PC and Atari STs can store around 720Kb on a double density disk, while a Mac manages to fit 800Kb on the same disk. There are also utilities around which cram a few extra per cent capacity on PC disks, by formatting them differently.

So while the disk manufacturer is not setting out to deliberately mislead us, there is no way that any computer is going to get 2Mb of data onto a 2Mb diskette.

Reading a bent disk

The other day, the editor of one of our sister publications brought me a 5.25-inch floppy disk and asked me if I could read it. This would not normally be a problem, except that it looked as if it had been rolled up and stuffed into a mail tube. The outer jacket looked like a write-off, but the disk itself didn't look too bad — there were only a couple of tiny creases in it.

So how do you recover data from such a disk? For starters, you need to want to recover the data. That is not intended to sound facetious — you already have one dud disk, and the recovery process involves destroying another one. So the data you're trying to recover has to be worth at least as much as a floppy disk.

Take an unused disk, and carefully prise the jacket open along the top edge, (furthest from the head access hole), and remove the disk from within the jacket. Cut the remaining flap off so that it doesn't scratch any disks as they are inserted.

Now turn your attention to the damaged disk, and remove it from its protective jacket in the same way. Bend it all the way back, so

that it is out of the way when you come to remove the disk. When removing the disk from the jacket, try to avoid touching the active surface area with your fingers. This will just add to the problems which you already have.

Once the disk has been removed from the damaged jacket, slip it into the good jacket from the 'sacrificial' disk, taking care to ensure that it is inserted the right way up. Now, put this disk into a disk drive and try to read it. If it is readable, quickly copy the data to a hard disk or another floppy. If you get errors reading some sectors, it's time to start using any file recovery utilities you have at your disposal, such as the Norton Utilities, or the Mace programs.

The trick is to get as much data as possible off the disk as soon as you can. If the disk is creased, the magnetic particles will quickly wear away from these areas, rendering that data unreadable. Also, this dust will build up on the heads, so it is a good idea to clean the heads after any operation such as this.

Hidden Windows

Windows 3.0 has two text files which contain important information telling Windows about the system on which it is running, and the applications which run under it — system.ini and win.ini. Although these files are created automatically when Windows is installed, and updated when new software or drivers are installed, there is often a need to modify them manually, such as when you remove an application from your hard disk — how many applications have you seen with an 'uninstall' option, which cleans up after a package is removed from

the disk? The system's autoexec.bat and config.sys files are two others which are often changed for various reasons.

The Windows Notepad is one way of doing this — double-click on Notepad, and then load the appropriate file into it for editing, or create an icon in a program group which fires up Notepad with the file already open. However, the first of these can be rather tedious, and you may not want to clutter up program groups with an additional icon for each file.

Windows actually has a special utility for editing these files hidden in the 'system' subdirectory, called sysedit.exe. Firing this up brings up the program with all four above-mentioned files in separate windows, ready for editing. Of course, any changes made will not take effect until Windows is restarted, or the system rebooted, as appropriate. To make sysedit appear as an icon in Program Manager, start File Manager, and re-size or move it so that you can see the desired program group in Program Manager underneath it. Then locate the sysedit.exe file, and simply drag its icon over the appropriate program group, and it will appear in Program Manager.

By the way, remember the old MS-DOS Executive front-end from Windows 1.x and 2.x? Well it is still there with Windows 3.0 and still called msdos.exe. For those times when you want to do a bit of file manipulation, but find all those windows in File Manager a bit cumbersome, the old MS-DOS Executive can still be useful.

Windows startup

Is it possible to disable the startup screen in Windows 3.0? It takes a good few seconds, which I am sure can be put to

better use than just displaying a pretty screen.

R. Bylart

Sure, it is possible to disable the Windows startup screen, by typing something on the command line, such as 'win :'. However, this will speed up Windows' startup little, if at all. Next time you start up Windows, watch the hard disk light — it doesn't stop flickering until the Program Manager window appears. Windows is using that time to load itself into memory, create temporary files and possibly a swap file. All this takes time, whether the screen is displaying anything or not.

Presumably the programmers at Microsoft thought users would like to see some feedback to indicate that Windows is actually starting up and hasn't crashed. Personally, I happen to like the opening screen for that very reason, (I probably crash my machine a lot more than most), and I like the colour scheme.

So get rid of the screen if you will, but you'll just end up staring at a blank screen instead of a startup message.

No graphics

I have a graphic display problem — in reverse! Is there any way to disable the graphic capabilities of IBM XT machines, most equipped with CGA cards, and one or two with Hercules-compatible cards. The students spend an awful lot of time playing games on these computers instead of (naturally) working on their assignments. What is worse, by borrowing and lending copies of games, the students continually infect the computers with viruses.

Most of the students' computer work consists of spreadsheet, word processing,

and database use, with only a limited need for graphics. Without graphic capability, the computers would become unusable for games and, hopefully, less prone to virus infection. One or two could be kept for graphic work, (and games when free). Can you suggest any way to disable the graphics?

G. Vatasan

That's an odd one all right! I'm not sure of any simple modification to CGA or Hercules cards which would disable the graphics modes, and a software solution may not be possible either, because of the way in which games tend to access system resources (such as the screen), directly, to maximise speed.

The only solution I can think of would be to network the PCs to a common file server, and equipping the network cards with auto-boot ROMs, so that they boot from the network, and not the floppy drives. By providing a user shell which restricts users to using those applications which you deem to be relevant to their studies, the playing of games would be prevented. You could still allow students to use the floppy drives to keep their assignments on, but not allow programs to be run from them.

If anybody out there does know of a simple way to disable the graphics modes on CGA and Hercules cards, drop us a line, and we'll pass the information.

RAM disks and Windows

I recently upgraded from my aging genuine IBM AT to a 25MHz '386 clone with 4Mb of RAM, due to the increasing demands of modern software on processing power, especially Windows and its multi-tasking ability. However, I am not sure

how to best utilise this amount of memory. Should I leave it to Windows to set up for me, or would I be better off using some of it for (say) a RAM disk or print spooler? Also, I had a 2Mb EMS board in the AT, and I don't know whether I should leave it there or transfer it to the '386.

B. Collins

Many people in the past bought high-performance machines with lots of memory, without realising that DOS wasn't capable of using the extra memory unless specific steps were taken to do so, such as implementing a RAM disk or disk cache. Windows 3 has changed this situation somewhat, with its ability to sense the memory configuration of the target machine and configure itself to make what it thinks to be the best use of the memory.

For most applications, Windows can be safely trusted to make the best use of available memory in '286 and '386 machines. If memory permits, it will install its own disk cache (`smartdrv.sys`), which uses some of the available memory to store recently-stored disk tracks. The beauty of `smartdrv` is that it is designed to work with Windows, so that Windows can dynamically re-size the cache as it needs memory, or has memory available for use.

RAM drives are rarely the best way to manage memory in a system — usually a better alternative is to use a disk cache. With a RAM drive, you have to decide which files to put in the RAM drive, and often you end up having to put files in the RAM disk which aren't accessed very often, simply because they need to be in the same directory as other files which are frequently accessed. Also, any data files stored in the RAM disk need to be copied back to the hard disk before rebooting. If the machine crashes before you get a chance to do this, it's too bad.

A disk cache, such as `smartdrv`, on the other hand, has some level of 'intelligence', which allows it to

determine what to put in memory, and what to forget about. Generally speaking, those parts of the disk which are accessed most often are those most likely to be found in the cache.

Be careful when comparing the speed of a machine with a RAM disk against a hard disk with a cache, as there are times when a RAM disk will appear faster, although this may not always be the case. For example, when launching an application from a RAM disk, it will almost always out-perform a cached hard disk. The reason for this is that launching an application typically means opening several files which haven't been used before, so they aren't in the cache. A RAM disk, which has the application files in it, will be able to make the files accessible to the operating system much faster than a cache, which first has to read the files in from disk.

However, when the program needs to access a single file (or a few files), several times in quick succession, the cache will rapidly approach the performance of a RAM disk, as more and more disk requests can be handled by the cache without reading the disk at all. The percentage of disk accesses which are handled entirely by the cache is called the cache 'hit rate'. The bigger the cache, the smaller the area of the disk which is accessed frequently, and the better the caching algorithm, the higher the hit rate.

So although you will never achieve a 100 per cent hit rate in a real-world application, the loss of performance compared with a RAM disk (which will always have a 100 per cent hit rate) has to be traded off against the extra time and effort needed to determine what should be in the RAM disk, and actually copying it there, and perhaps copying it back afterwards.

Also, because a disk cache caches disk sectors or tracks, only the most frequently-accessed parts of files need to be cached — you can't

store part of a file in a RAM disk, and leave the rest on the hard disk. Also, there is the security of knowing that anything written to the cache will eventually find its way onto the spinning metal of a real disk drive — not so with a RAM disk.

Don't be tempted to use a RAM disk as a swap disk either; the memory would be better used as system RAM by Windows or any other multitasker you are running. If the RAM used for the disk was directly available to the system, there would have been no need to swap to disk in the first place. If you have more than one hard disk in the system, put the swap file on the fastest one, since it is acting as a virtual RAM, and the faster the swap disk, the higher the overall performance of the system when things start to get busy.

So, in summary, unless you have a specific need for a RAM disk, it is almost always better to allocate the memory which you would use for the RAM disk, to a disk cache instead.

As for your EMS card in the AT, I'd be inclined to leave it where it is. You could conceivably put it in the '386, and use it as a RAM disk or disk cache, but if you really wanted extra memory in the '386, you would be better off investing in a proper 32-bit memory card to suit your motherboard — the memory accesses will be much faster, and you can emulate EMS if required on a '386 with something like QEMM (from Sourceware, (02) 427 7999), and Windows can also emulate EMS for applications running under it.

Also, if you use the board in the '386 with its EMS driver, there is a good chance that it won't work with Windows in Enhanced mode, which is rather fussy about the company it keeps. All things considered, trying to get the EMS card going in the '386 is probably more trouble than it's worth, but don't let me stop you trying it out, if you're keen. If it was me, I would probably have a play with it, if only to satisfy my curiosity.

Tech Tip of the Month

I HAVE A Toshiba T1000SE notebook computer, which has just the one 3.5-inch floppy drive. Because the MS-DOS Backup and Restore utilities are not suitable for backing up diskettes on a single drive, I developed a small batch file which makes backing up simple, and also keeps a record of what was done and when it was done.

Backprod.bat takes advantage of the 328Kb RAM drive (drive D:) which the Toshiba maintains with its battery, even when switched off. This eliminates any problems that would arise with the batch file being placed on the drive which is being backed-up.

The batch file expects one parameter — the volume label on the diskette to be backed up. This is for documentation reasons, as the backup will proceed whether the volume label on the diskette matched the command-line parameter or not. By logging the backup to a log file (also on drive D:), I can always check which volume was last backed up, and when. It also reminds that backups haven't been done for a while.

The batch file starts by displaying a banner, and then checking for a parameter. If the parameter is 'help', or no parameter is supplied, control is transferred to the 'helpinfo' label, where a help message is displayed. Otherwise, the parameter is appended to backprod.log, then by the actual volume label on the diskette, and then the backup begins using the diskcopy utility. I chose diskcopy as it handles disk swaps in the one drive happily.

When the diskcopy program is finished, the success or failure of the backup is checked using the errorlevel test, and an appropriate message displayed on the screen, and recorded in the log file. The time and date are also recorded in the log file for future reference.

```
@echo off
echo _____
echo          backup procedure for diskettes
echo _____
;check for parameter
if %1 == "goto helpinfo
if %1 == 'help' goto helpinfo
if %1 == 'HELP' goto helpinfo
;perform backup and log to log file
echo Backup of %1 diskette begins>>backprod.log
vol a:
vol a:>>backprod.log
diskcopy a:a:
;check for success or failure
if not errorlevel 0 goto error
echo Backup of %1 completed successfully>>backprod.log
goto endjob
:helpinfo
echo          Enter BACKPROD volumename
echo          to run this backup
echo          .....
echo          note that it logs to BACKPROD.LOG
echo          .....
goto endall
:endjob
echo          **** End of BACKPROD job ****
now >>backprod.log
echo          **** End of BACKPROD job ****>>back prod.log
:endall
echo on
```

I have found this batch file very useful in overcoming the only problem with my precious Toshiba.

Steve Clark

Notes: The 'now' command used in the batch file is a utility supplied with the Toshiba in its ROM version of DOS, which prints the current date and time on the screen. If you use this batch file on

another type of laptop, you could get away with using the DOS time and date commands. If you create a file with just a carriage return in it, and re-direct this as input to the time and date commands, you won't have to hit the Enter key each time.

The only disadvantage is that the 'Enter new date/time' messages will appear in the log file. Norton's time mark (tm) utility would be a better solution.

Static protection

With RAM chips costing but a fraction of what they did a year or so ago, I decided it was time to upgrade the memory in my AT from the supplied 1Mb, up to 4Mb. In my machine, this simply involves (according to the manual) setting a few DIP switches, and replacing the 36 RAM chips with bigger (1 megabyte) numbers. However, after buying the chips, I recalled reading an article which warned about the possibility of static damaging the chips when replacing them. However, the retailer who sold me the chips made no mention of it, and I can't find any reference to it in the manual for my computer either. How important is it to protect against static, and more importantly, how does one go about it?

M. Bannister

Static damage to chips is one of those things which is hard to pin down. Many people (including me), have gotten away with handling chips without taking any special precautions, and have caused them no apparent damage. On the other hand, there is no doubt that static charges can destroy the tiny junctions in integrated circuits, but if a chip does turn out to be dead, there is usually no way of finding out whether the cause was static, or something else.

Static is thus a little like AIDS — you might get away without precautions for a while, but eventually the odds will catch up with you, and you'll zap something. In the case of a RAM chip, the problem can be particularly insidious, as it is often difficult to tell which chip is at fault, and the problem can often appear to be intermittent in nature, making tracing the fault even more tedious.

Protection can take one of several forms, depending on the importance of the protection, and the level of static which could be

present in the environment. The simplest way to prevent damage is to simply touch the metal chassis of the computer before plugging the chips in. This will equalise any static potential between you and the computer which you may have built up beforehand.

However, if you are standing on nylon carpet for example, particularly in dry weather, simply shuffling one's feet around in the course of installing the chips could build up enough charge to damage a chip or two.

A better solution is to use a conductive wrist strap, as sold by most good computer parts outlets. One end of the strap is wrapped around your wrist, while the other end is attached to the computer which you are working on. This discharges existing static as before, but because it stays connected, any charge which is generated during the course of the operation is immediately drained away before it reaches a dangerous level.

Electronics buffs usually have a number of jumper cables lying around, which consist of a short piece of wire with an alligator clip at each end. One of these clipped to the computer chassis and a metal watch band (which you're wearing, of course), will achieve the same result.

Also, the forgoing applies to add-on boards, as well as single chips. Although external circuitry on a circuit board provides some level of protection to the devices on the board, there is still a chance of damaging chips on the board with static, and because they are typically soldered directly onto the board, replacement is not a simple task.

When handling boards, it is good practice to avoid touching the connector pins or the tracks on the board — handle the board by the edges as far as possible, and if you do have to touch the metal tracks on the board for some reason,

consider grounding yourself first, as described above. SIMM RAM modules can be considered to be bare chips for the purposes of static protection. Although SIMMs consist of a few small chips soldered to a small board, the latter provides no protection as far as static is concerned until plugged into the computer.

dBase and DR-DOS

I recently installed dBase IV, version 1.1, into my Datamini 286 AT, which is attached to a Star NX-1000 printer, and running the DR-DOS 3.3 operating system. Every time I use the printer with any of the dBase print commands, the computer seizes, and I have to reboot to start the program again. As a result, I cannot use dBase IV on this computer. The printer driver has been properly installed.

An Ashton-Tate technician informs me that this is because DR-DOS 3.3 and dBase IV are incompatible. In what way can I solve this problem?

M. Mohantal

Unfortunately, I cannot think of any way to solve your problem other than changing your operating system. Just as the first attempts at building IBM PC-compatible computers had all sorts of incompatibilities, the same seems to be the case with the DR-DOS operating system. With release 5.0, Digital Research has addressed a lot of the compatibility issues, but it is still far from perfect. I recently tried out a copy, and couldn't get some device drivers to work with it. I didn't have the opportunity to try it out with dBase IV, so I can't comment on that combination specifically. My recommendation would be to get a copy of genuine Microsoft MS-DOS, and change your system over to that.

April 1991

More IRQeries

In 'Tech Tips', December 1990, you comment that the interrupt line IRQ2 is hardly ever used in ATs and later indicate that it is available. However, the instruction booklet for the Inport bus interface of my Microsoft mouse specifically states not to use this option for AT computers.

My mouse was originally installed on an XT for which the manual excludes the use of IRQ5. Interrupt lines 3 and 4 are reserved for COM ports, and the manual states that pin pair 2 is the default setting, and is covered by the jumper clip. Your article states that IRQ2 is reserved for the hard disk on XT machines!

When I changed to an AT computer, the Inport bus interface was moved over from the XT without checking the jumper block was correctly set. The mouse functioned with no apparent problems on either machine, so is there any reason why I should change the jumper to IRQ5?

I have had periods of intermittent crashing of the AT along with a message about parity errors. This was one of those embarrassing faults that did not occur when I took the computer back to be serviced. Eventually the RAM chips were replaced, and so far the fault has not recurred. Is it possible that the parity error problem is related to my mouse being connected through IRQ2?

K. Glasziou

Firstly let me apologise for an error in the article to which you refer — the hard disk interrupt line on XT machines is indeed IRQ5, not IRQ2 as originally stated. IRQ2 is unused (or reserved, in IBM's words), as you found when using it for your mouse.

There is a fair amount of confusion amongst users (quite experienced ones, too), regarding the use of the IRQ2 line in AT machines, which often leads to people leaving a perfectly usable interrupt line unused. Strictly speaking, the IRQ2 line in the AT is

connected to a slave interrupt controller, which gives the AT seven more interrupt lines than the XT — the second interrupt controller chip handles eight lines, but we lose one line — the IRQ2 line — on the primary controller, leaving us with a net gain of seven.

One of those new lines, IRQ14, is the new hard disk interrupt, leaving IRQ5 free for the second printer port. Two of the other additional lines (IRQs 8 and 13), are used by the real time clock and maths co-processor, leaving five spare lines on the second controller. Four of these lines are connected to the extended expansion slots, and can be used by peripherals designed to take advantage of this. The remaining line, IRQ9, is connected to the pin on the expansion bus that is used by IRQ2 on XT machines. So, any card which has a jumper to select IRQ2 will in fact be using IRQ9 when plugged into an AT.

Software has no problems with this, as the AT's BIOS sets the interrupt vectors so that the vector for IRQ9 is the same as that for IRQ2 on an XT. So, as far as the software is concerned, the device is using IRQ2, be it in an XT or an AT.

So, by selecting IRQ2 on your bus mouse, you are in fact using IRQ9, but the software thinks it is IRQ2. I cannot understand why Microsoft says not to use IRQ2 in ATs, except possibly that the writer of the manual saw that IRQ2 was used by the second interrupt controller, and was unaware of the IRQ9 redirection. This is a common misconception amongst users and developers alike.

I have installed Microsoft bus mice on several computers, and all of those have used IRQ2. None of these machines has exhibited any interrupt-related problems. In any case, if the Microsoft Mouse driver cannot find what interrupt line the mouse is using, it will deliver an error message and not install itself.

The RAM parity problem which you met is not related to the interrupt lines at all, but is most probably due to a below-spec RAM chip, or one that is not rated for the speed of the machine. Since you don't state the speed of the machine or that of the RAM chips, it is impossible to determine whether there was a faulty chip, or whether the chips were being called upon to work faster than they were designed to. However, since the chips have already been replaced, and the problem solved, the problem is now academic.

DOS + Xenix

I trust you may be able to answer a query I have regarding the running of both DOS 3.3 and Xenix on the same PC. I operate an NEC Powermate 386/20 with 4Mb of RAM, a 140Mb hard drive, Intel 80387 maths co-processor, twin floppies, enhanced VGA and a 101-key keyboard.

What I would like to do is to alternate between DOS and Xenix as efficiently as possible. I have special programs that are written in C and utilise multiple screens, plus I am looking to test drive products like 20/20, all of which operate under Xenix. As well, I have a multitude of applications which operate under DOS. I would appreciate any suggestions that you can offer which will allow me to alternate between the two operating systems rather than having to go through the fdisk program each time.

The 'Tech Tips' section of your magazine has been a source of invaluable assistance, which proves that readers like me can and are becoming increasingly frustrated with the lack of support and understandable manuals for both hardware and software products.

M. Gill

We're glad that 'Tech Tips' is fulfilling a need for PC users — a

need that perhaps should not be there in the first place. I cannot agree more with you on the subject of proper support and documentation for both hardware and software. It seems that, in their quest to make manuals easy to follow for novice users (a goal which, by and large, most have missed almost entirely), manufacturers are neglecting more advanced users, who want to do something a little unusual.

Back in the 'good old days' both software and hardware manuals tended to give a great deal of information, so that the user had the best chance of getting the product to work with his or her existing system. This usually meant that for a given user, most of the information presented was unnecessary and obscured the relevant material.

Modern manuals, on the other hand, have gone to the other extreme. Rather than putting material that is of no interest to most people in an appendix, where it is out of the way, yet still accessible to those who are interested, many seem to omit critical pieces of information, leaving you to ring their tech support department (which is usually in the States), and try to find someone there who even understands your problem.

Unfortunately, with the sheer amount of hardware and software available on the market, it is impossible for a column such as this to be able to provide immediate answers to all people's problems. What we aim to do is to provide a forum where users can exchange hints and tips with other users using similar products

Anyway, that's a bit of a side track. Your problem is one that a lot of people will never encounter, simply because they run DOS on their PC, and nothing else. Even those who wish to run OS/2 do not encounter the same problem, as both DOS and OS/2 use the same hard disk structure.

However, as you well know, Xenix is a different matter entirely, using its own disk directory structure, which is completely incompatible with DOS. To have DOS and Xenix on the same hard disk, you need to partition the hard disk, and install each operating system in its own partition, as you have done.

Modern manuals, on the other hand, have gone to the opposite extreme.

When the machine is powered up, and the BIOS finds a hard disk, it looks at the partition table to see how many partitions are defined, and which partition it should boot from. The partition from which the system is to be booted is known as the active partition, and for most systems, this is the primary DOS partition. Since most systems only have at most one primary DOS partition, and possibly one extended DOS partition, which is not bootable anyway, most people don't need to change the active partition.

However, with two operating systems on the same hard disk, there needs to be some way for the BIOS to know which partition of the hard disk to boot from. To this end, the partition table has an entry for each partition, which specifies whether it is active or not. Although it is possible to have more than one partition marked as active, it is likely to confuse the BIOS, so this shouldn't occur.

DOS uses its fdisk program to edit the partition table, but it is limited in the operations that it can perform. It can create and delete DOS partitions, and can make partitions active or inactive. However, it cannot create or remove partitions for other operating systems, such as Novell or Xenix. So if you have a non-DOS partition that you want to delete, you can't use fdisk to do this. Not that this is a problem in your case, since you already have the disk partitioned as you need it.

Fortunately, Norton Utilities includes, in its Disk Editor, the ability to edit the partition table. Norton recognises most types of disk partitions, and enables modification of any parameter in the partition table. This editor is extremely powerful, and it is easy to trash the partition table, rendering the entire disk useless.

However, changing the active partition with Disk Editor is really no more convenient than using fdisk. However, there is another utility in the package called Disk Tools, which can create and restore 'recovery' diskettes, which back up, amongst other things, the partition table. This can be used to switch rapidly between two different copies of the partition table.

The first step is to take two blank floppy disks and copy the system files (and command.com) onto them, using the sys command, or format with the '/s' parameter. Also copy disktool.exe to each floppy and create an autoexec.bat file on each disk with a single command — disktool — in it

Using fdisk, set up the partition table, so that the DOS partition is active and the Xenix partition inactive. Now run disktool, and select the 'create rescue disk' option, saving the files to one of the floppies created before. Label this disk 'DOS'.

Now, run fdisk again, and make the Xenix partition the active one,

Tech Tip of the Month F11 and F12

IN THE December, 1990 Tech Tips, I discussed an easy way to re-define the PC's function keys using the ansi.sys console driver. However, this only covered the Function keys F1 through F10, leaving the two additional keys on the extended keyboard undefined.

A frequent correspondent with Tech Tips has found a solution —

I solved the function key problem with a small shareware TSR called F11F12.COM, which takes up 688 bytes of RAM. It came on 'Disk 710, Utilities for the AT', from Shareware Marketing, Beer, EX12 3HW, England.

I recently acquired a Toshiba T1200XE, with MS-DOS 4.0, which includes an updated ansi.sys driver. Putting

DEVICE=ANSI.SYS /X

in my config.sys file was sufficient. The '386 mentioned before has DOS 3.3 and this rejected the Toshiba ansi.sys.

Using the 4DOS command processor, I now have 40 single-stroke function keys to do most of my DOS work plus a spare for standard 4DOS functions. (for example, F1 is DOS help).

However, there seems to be a limit to the number of key strokes that can be stored by the ansi key re-definition method, so I used 4DOS aliases to reduce each command to a single character in the ASCII 129 to 255 range, and put a template over the function keys so I know what the hieroglyphics mean! In so doing, I found another problem — a number of the ASCII codes such as 'â', 'ä', and 'a' are all treated as 'a' by 4DOS aliases; something to do with the collation sequence, no doubt.

E.S. Webber

Mr. Webber also provided a copy of the on-disk documentation for the utility, which goes on to say that the program also allows the shifted F11 and F12 keys to be read by ansi.sys. Note that the program does not do that for the unshifted keys — it actually puts the characters which you define into the keyboard buffer. If your local BBS does not have F11F12, the program's author is:

Chris Rabinowitz
51 East Rogues Path
Huntington Station, NY 11746,
USA.

and then re-boot from the second floppy disk. Again run disktool, and save the 'rescue' disk information to the second disk. This will become your Xenix disk.

Now whenever you want to change the operating system, boot up from the appropriate floppy, and then follow the menus in the disktool program, to restore the partition table information. It is wise to restore this information only, and not the boot record or CMOS RAM information as well, so de-select these options before actually selecting 'OK' on the last screen. Once the information has

been restored, open the disk drive door, and select 'reboot'.

I'll admit it's still a bit tedious, but it should be a bit simpler than using fdisk constantly. Another thing you might like to try is using a keyboard 'buffer-stuffer' to automate the keystrokes in the disktool program, so that the entire process is automated, except for opening the disk drive door when the re-boot commences.

The utilities 'Keyfake' or 'Putkey', which are available from many BBSs, would automate the process entirely, presenting the appropriate keystrokes to diskedit

as it requires them. To work out what keystrokes need to be entered, go through the restore process manually, and note down every key you press. Then consult the documentation of the utility to see how to specify these exact keys to the program.

I haven't tried this last part, but I can see no reason why it shouldn't work. The only possible problem that I envisage is that you may end up with more keystrokes than fit in the keyboard buffer. In this case, there are a large number of utilities on BBSs which allow you to extend the size of the keyboard buffer. None of this will effect the normal operation of the machine, since the machine is re-booted (under either DOS or Xenix), after restoring the appropriate partition table, so any changes made to the keyboard buffer are only temporary.

Hopefully, this gives you an idea of what can be done with commercially available and public-domain software. Of course, the best solution would be to write a little assembler program which directly modifies the partition table's 'active partition' entries. At least, using the above method, you know that you are restoring a valid partition table that was created by fdisk.

More on more COM ports

One of the many limitations of the hardware configuration of the PC is the lack of interrupt lines available for use by add-on boards. The AT has improved this somewhat, by almost doubling the total number of lines, but a lot of hardware and software doesn't support these extra lines. However, with some of the standard devices in the machine now using the new interrupt lines, there are more

options available for the low-order lines. Since we're already talking about interrupts, we might as well stay on the subject for a while.

One of the most frequently encountered problems arises when one wants to add more than the standard two serial ports to a machine, for whatever reason. Since most software using a serial port is interrupt-driven, and written with the assumption that it has exclusive use of a given interrupt line, each serial port needs to have its own interrupt line if they are to be used simultaneously.

With a mouse now virtually standard equipment, and machines with a dedicated mouse port still by far in the minority, a machine with only two serial ports has one permanently occupied by the mouse. Add a modem, and you have none left for anything you want to do in the future. My machine at home currently has a mouse, modem, packet-radio controller, and a null-modem connection to my laptop. I can use (and have), all these ports at once (really, only to see if it could be done), without any conflicts. To do this, I installed a second serial card, and set the addresses for the new ports to the standard COM3 and COM4 addresses. I also hacked the board a bit, to change the interrupt lines.

The two standard serial ports use interrupt lines 4 and 3, for COM1 and COM2, respectively. Most implementations of COM3 and COM4 (including my card), use the same pair of lines, sharing them with the two existing ports. Since DOS's serial port routines are polled rather than interrupt-driven, DOS happily supports four ports, since it doesn't use the interrupt lines at all. However, the DOS routines are also slow, so no programmers use them, writing serial service routines instead.

This is where the problem with the sharing of the interrupt lines

arises, since the two drivers monitoring the same interrupt line don't know about each other's existence. There is also a hardware problem with sharing the interrupt lines, due to IBM's choice of active-high levels, rather than active-low levels, which would allow the use of open-collector drivers on the cards, and allow interrupt sharing, at least at a hardware level.

... it means a bit of surgery with a scalpel and — gasp! — a soldering iron.

So, the only real solution is to give the new COM ports their own interrupt lines, and unless the card has jumpers to allow you to do this, it means a bit of surgery with a scalpel and — gasp! — a soldering iron.

Before commencing, the first thing to decide is which interrupt lines to use. IRQs 3 and 4 are out of the question, since they are used by the first two serial ports. Ditto IRQ6, which is used by the floppy drive, and IRQs 0 and 1, which are used on the system board for the keyboard and timer. Also, IRQ5 is used by the hard disk controller in XT machines, which leaves IRQs 2 and 7, and also 5 in ATs.

Now, IRQ7 is reserved for the hard disk, but in practice, it is rarely used, except by some print caching programs. The same also applies to IRQ5 on ATs, which is reserved for LPT2. (Why does the AT need separate lines for its two printer ports when XTs don't? Because it doesn't, in the same way as LPT1 doesn't use an interrupt except in special cases.) IRQ2 (really IRQ9 on

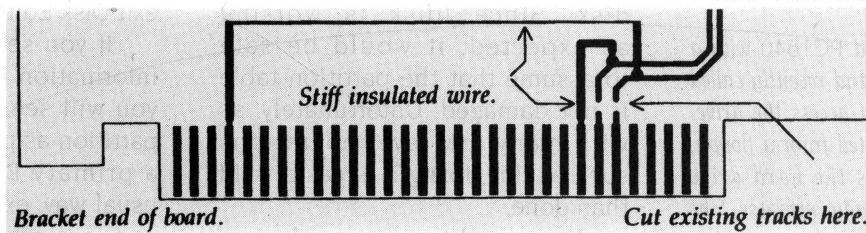
ATs — see previous explanation), is not used on either machine.

I eventually settled on IRQ2 for COM3, and IRQ5 for COM4 in the Syncomp Mega 386i which I'm using as my main testbed at the moment. The I/O addresses for these ports stay at the standard settings for COM3 and COM4, and any program which doesn't use the interrupt lines (such as LapLink III), will still be able to use them if desired.

Now, down to the details of modifying the card. First of all, you have to select a suitable card for modification. You have the choice of modifying the I/O card already in your machine, and buying another to be used for your first two ports, or you can modify the new card. First of all, the chosen card needs to have jumpers to select the ports as COM3 and COM4. If you can't do that, you won't get past first base.

Ideally, at least one of the cards should have a socketed UART chip (the big ones with 40 pins), which will allow you to upgrade it at some future time for a high-speed modem, should you buy one. Cards with a single surface-mount chip which contain both UARTs are all right, but you won't be able to change the UARTs later on if you need to. The card should also have a complete set of 'fingers' along the expansion bus edge-connector. If there are any missing ones, find another card.

Now for the bit with the scalpel. Turn the card to be modified over, so that the side of the board with the chips on it is facing down, and you are looking at the side with all the solder on it. Starting at the end of the edge connector furthest from the metal bracket (if the card is a 16-bit one, ignore the 'extra' AT contacts, and start with the contacts on the larger part of the edge connector, common to both XT and AT machines), and count along to the seventh and eighth contacts.



Follow this diagram carefully to determine which tracks to cut, and where to solder the wires. The board is viewed from the read (the side without the components on it).

These are the IRQ3 and IRQ4 lines, and there will be thin traces leading away from these pins, leading off into the depths of the board somewhere. Carefully with a scalpel, cut these tracks, as close to the edge connector as possible. A magnifying glass is useful here to make sure that no slivers of metal remain to maintain a connection.

The next pin along is IRQ5. If there is a parallel port on the card as well, there may be a trace leading away from this one also. If IRQ5 is one of the lines you intend to use, cut this trace too. Skip the next pin on the connector, and look at the one following, (IRQ7). Again, if you intend to use this line, and it already has a trace leading from it, cut it. IRQ2 (if you are using it), won't be connected unless you have something strange on the card besides I/O ports.

Now comes the time to warm up the soldering iron. If you don't have some soldering experience on tightly packed circuit boards, it is best to get somebody who has worked on this sort of gear before, to do it for you.

For my set-up, I wanted the track which went to IRQ3 (COM4), to go to IRQ5. To this end, I scraped about a quarter of an inch of the solder mask (that green stuff covering the copper tracks on the board), from the severed track which used to go to IRQ3 on the

edge connector, and soldered a short piece of insulated stiff wire to it. The other end was soldered to the finger on the edge connector for IRQ5. Make sure that the solder does not run down the track of the edge connector, or it will make inserting the card into an expansion slot difficult.

Repeat the same procedure for COM3, which has to go to the connection for IRQ2 on the edge connector. (See the accompanying diagram for more details.) Once you have finished, press the wires flat up against the surface of the board, so they do not snag on other boards, when you are inserting the new board. Before plugging the card in, make sure that you set the addresses of the ports to COM3 and COM4, according to the manual for the card.

That just about completes the process. Now all you have to do is to get your software to recognise the new interrupt lines. Telix makes this easy with its comm port setup option, accessible from the configuration menu (Alt-O), and changes to appropriate entries in the displayed table to reflect the new interrupt lines. Now, whenever you select either of the second two COM ports, Telix will know what interrupt line they are using. Note that Telix doesn't recognise the extended interrupts on ATs so you have to use IRQ7 or lower with

Telix. In addition, it is oblivious to the IRQ2 redirection mentioned above, so that it sees the XT IRQ2 line (IRQ9 on ATs), as IRQ2.

The choice of what you use each COM port for is largely a matter of the flexibility of the software which is driving each port. On the Syncomp here, I have the mouse on COM1, since the driver requires it to be on COM1 or COM2. COM2 is used for a null modem connection to my laptop, and as a general I/O port at other times. COM3 and 4 are used for my modem and packet radio set-up, since Telix can handle virtually any address/interrupt combination.

Windows 3.0 in Enhanced Mode is also supposed to allow setting of the COM port base address and interrupt lines, but I haven't quite got that one to work yet. When I do I'll let you know.

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on.

The best *Tech Tip* published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is Tech Tips, *Your Computer*, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

May 1991

Partitioning damaged hard disk

A client has an Amstrad PC1640 with a 30Mb RLL hard disk, and recently called me because she couldn't access the drive. When the system is booted from a floppy, any attempt to access the hard drive results in a 'drive not ready' message. She had run Norton's speed disk (version unknown), and apparently there was some disturbance in the power during that operation. I couldn't access the drive at all, so I decided that a low-level format would be the only way to get it going properly — fortunately she has a backup. I have done several without any problems, using G=C800:5 in debug. The low-level format seemed to work — the message 'format successful' returned each time. I then used fdisk to make a DOS partition, but the high-level format failed after trying for a couple of seconds. I presume that this is something to do with the RLL controller. I've asked everyone I can think of, with no solutions offered. Other than going to Amstrad for a (possibly expensive) service, is there anything else that can be done to resurrect the drive?

P. Dalman

It's good to see a problem from somebody who has backed-up their drive, for a change! It certainly makes the process of fixing the problem a lot simpler if you can just ignore the data on the disk and start from scratch.

What seems to have happened is that a bad sector has developed on the disk where the important, non-movable, areas of the DOS partition need to go — such as the boot record and file allocation table (FAT). While DOS can lock out bad sectors in a partition, it stores this information in the FAT, so that the FAT, and those areas of the disk preceding it, cannot work around bad sectors.

So what needs to be done is to start the DOS partition not at the very beginning of the disk, but a

cylinder or two in from the start. This will put these critical bits in a (hopefully undamaged) part of the disk. Since fdisk is working as expected, it would be safe to assume that the partition table is not damaged. Unfortunately, as with many low-level operations such as this, fixing it is easier said than done.

Your need to create a dummy partition of one cylinder, and then create a primary DOS partition starting immediately following it. Unfortunately, this cannot be done entirely with fdisk, but as you mention that Speed Disk was running at the time of the failure, you obviously have a copy of the Norton Utilities. Either the Version 5.0, or Advanced Edition of versions 4.0 or 4.5, have a suitable partition table editor, built into the Disk Editor program in version 5.0, or NU in 4.0 and 4.5. Although I am going to describe the steps relevant to version 5, there are equivalent operations which can be performed using the earlier versions to achieve the same result. Consult the manual for the utilities if you are unsure how to access the required menus.

The first step is to delete any existing partitions on the disk, and create a primary DOS partition of one cylinder in size. Fdisk will insist on re-booting the machine at this point, after which you should fire up Disk Editor (or whatever it is that you are using). Select the first hard disk as the object to modify, and then select 'partition table'. The partition table should come up as a table with four entries, a DOS partition, and three unused ones. Make sure that the editor is in partition table mode, by pressing F6, if necessary.

Move the cursor to the name of the first partition (probably DOS-12, as it is a small partition), and hit the space bar to change it to something else — it doesn't matter what, as long as it is not any

type of DOS partition, or 'unused'. Xenix would be a suitable choice. Write the changed data to the disk, exit Disk Editor, and re-start fdisk.

If you select 'display partition information' from the fdisk menu, you will see a one cylinder Xenix partition as partition 1. Now create a primary DOS partition in the usual way, except that there is now one less cylinder of space — this space is lost forever, unless you reverse this procedure. Before quitting fdisk, make sure that the primary DOS partition is selected as the active partition.

Provided that the only bad sector was on the first track of the drive, then format should work in the usual way, except that you will have lost 1/n of the capacity of the drive, where n is the number of cylinders on the drive. The extra cylinder, which is marked as a Xenix partition, is never seen by DOS, and the BIOS won't try to boot from it, since the second partition is the active (bootable) partition.

By the way, I don't think that the fact that Speed Disk was running when the power failed was the direct cause of your problem. While any program writing to a disk when the power fails can cause damage to the data on the disk, this is not due to a physical problem on the disk, but a logical error where the file allocation tables do not match the contents of the files themselves. Such a problem would be cured when the disk was re-formatted, thus erasing all of the information in the FAT and root directory.

It seems to be that the disk suffered some physical damage, possibly due to the heads not being parked properly when the power failed, causing them to crash onto the surface, and damaging it. Hopefully this hasn't damaged the heads themselves, so the drive may still be usable by just ignoring the damaged cylinder.

MAY 1991

Constant Speed Modems

I recently upgraded my modem from a NetComm 1234SA, which has served me well, and purchased a new V.32 M5. While the initial outlay was not insignificant, I am a long way from the nearest capital city, where all of my favourite BBSs are located, and I expect that it will pay for itself in reduced telephone charges pretty quickly. While many of the BBSs that I use have V.32 modems, not all of them do, and so I am still forced to use 2400bps on some boards.

However, I am having difficulty with the speed settings of my new modem, and the communications program which I use (Telix). The modem manual says that I should set up my communications program and modem to run at a constant 19200bps, and use handshaking to keep the computer and modem in step with one another. But, whenever I do this, I can no longer connect with 2400bps bulletin boards. If I slow down the terminal speed of the computer to 2400bps, before dialling, then I connect alright, but I don't get any speed advantage from MNP, which one of the boards uses (I think they have a Dataplex 224). So, is there any way that I can get the modem to connect at 2400bps, without restricting the terminal speed to that figure, and throwing away any advantage of MNP — every little bit helps when you're calling STD. Do I need new software (I hope not, I like Telix, and being Shareware, it is cheap)?

J. Seeto

Your problem is a product of our collective desire for faster and more accurate communications, but fortunately there is a solution. As you probably know, when dialling a bulletin board, the answering modem 'auto-ranges', to determine the speed of the modem at your end. If you dial a 9600bps modem from another modem of that speed, they will connect without any trouble. However, you can also dial the same modem from a modem which can

only run at 2400bps, 1200bps, or even 300bps. The answering modem steps through its collection of operating speeds, sending the answer tone for each speed, until the remote modem responds — indicating that it has found the operating speed of the remote modem. This is auto-ranging.

If, however, you dial, say, a 2400bps modem from a 9600bps modem, they will never establish a connection. This is because the remote mode cannot generate the answer tone that the 9600bps modem is expecting — it will cycle through the speeds which it does know about — most likely 2400, 1200, and 300bps, but the calling modem is listening for a 9600bps carrier, nothing else.

This is because calling modems do not auto-range. They listen for a single carrier tone, and only respond when they hear the one that they are expecting. This is done to prevent the two modems ranging up and down independently through their respective ranges, and missing each other each time. It's a bit like arranging to meet somebody in a shopping centre, but not specifying where. You both arrive at the allotted time, but not knowing where you are to meet, the chances of both of you being in the same place at the same time are rather slim. Telecom won't let modems keep the phone line if they are not communicating, so eventually both modems must hang up, without ever having exchanged a single bit of data.

In order to determine which speed a modem is to call at, it will generally look at the line speed. If the dial command from the computer is sent at 300bps, then the modem will expect to connect at 300bps. So, if you dial a 2400bps modem from another 2400bps modem, but issue the dialling command at 300bps, then the modems will operate at the slower speed, even though they are both capable of running at eight times

that speed. This can be used to advantage if you are calling a BBS over a particularly bad line, and want to slow down the communications speed to minimise errors.

Your M5 works in the same way. If you issue a dialling command at 2400bps, it will either connect at that speed, if the remote modem is capable of that speed, or the connect will fail. There is no half-way. Ditto for 9600bps. If you send the dialling command at 19,200bps, the modem will still connect at 9600, since it can't run any faster than that — it will try to connect at its highest available speed. If a 2400bps modem answers, yours won't hear it, and will eventually stop trying to connect.

What you need to do is to force your modem to call at a slower speed (2400bps), even though you send the dial command at 19200bps. In NetComm modems, the ATB command is used for this purpose. If you type ATB8 before dialling the remote modem, then your M5 will call at 2400bps. When a 2400bps modem answers your call, the two will connect, since your modem is now, to all intents and purposes, a 2400bps modem. However, since your line speed is still 19,200bps, you gain the benefit of the extra throughput which MNP allows.

But, if you now call a 9600bps board, you will only connect at 2400bps, as your modem no longer recognises the V.32 answer tone. You need to put the modem back in its original mode, by typing ATB0.

This can get a bit tedious, so the best thing to do is to use Telix's multiple dialling strings to do the work for you. Set up two separate strings — one for 9600bps boards, and the other for 2400bps. Make the first dialling string ATB0D, and the second one ATB8D. Now when you dial a BBS, the modem will be kicked into the correct mode automatically before dialling, so that a connection can be established with ease, and still gain the extra throughput offered by MNP.

MAY 1991

Extended memory loss

Having read the letter in the July 1990 issue of 'Tech Tips', and your response, I quickly checked my May 1990 issue to do as you suggested. I read and re-read the article by Stewart Fist, and promptly did as his last paragraph instructed. I know nothing about the use of memory — be it extended, expanded, or whatever.

Well, that's not strictly true! I have come to know about a sort of memory not mentioned in his article, namely 'low memory'. It seems I have quite a large chunk of it.

Having just entered the world of IBM/MS-DOS computing (I spent the last 5 years playing with an Apple IIe), I was enjoyably amazed at the vast differences. The system I chose was an ALR '286/18MHz Powerflex Plus, with 1Mb RAM, with a variety of peripherals — floppies, hard disk, and printer. I also purchased Windows 2.11, Word for Windows, Paintbrush, and a mouse.

I began, with the help of an experienced computing friend, to arrange the best possible setup using the system I had purchased. It seemed I had, as far as memory goes, 640Kb base memory, and 384Kb extended. My friend seemed satisfied that this should be sufficient for my initial needs, and it was to a point.

Having set up the `autoexec.bat` `config.sys`, and `win.ini` files, I was away. Printing my sort of memos, letters and other text-oriented documents was a breeze. I was printing documents, which on the Apple IIe took several files, as a single file with a myriad of fonts, colours, and sizes. Lovely stuff!

Then I became rather daring — I started to incorporate a graphic or two. Just practice, mind, nothing way out. A simple use of the few TIFF files that came with Word. But that was where it all came unstuck, and I developed my use of that large block of memory known as low memory. I spoke to a guru or two, but they said my system was too small. I needed more memory chips. So I got them. Only, the dreaded 'low memory' screens still popped up.

Then I noticed a difference in the initial screen you see at startup. Instead of 640Kb base, and 384Kb of extended, the extended memory was no longer there. Instead, a new message indicated that I now had 64 pages, or 1024Kb expanded memory available. Now, that just didn't add up. I had started out with 1024Kb (640+384), and had added another 1024Kb, but now only had 1664Kb (640+1024). Where had my 384Kb of extended memory gone? Then it occurred to me. That must be the part showing up as low memory. So now I had 640Kb base, 1024Kb expanded, and 384Kb low memory; yes?

The guru or two spluttered and mentioned a few expletives, but were unable to explain where my 384Kb of memory had gone. A quick call to the ALR distributors in Brisbane gave little joy. The response was: 'You can't have extended and expanded memory together.'

Does Windows 2.11 make use of it? If not, would an upgrade to Windows 3.0 help? I have used `SmartDrive.sys` in the `config.sys` file, but get better results when using the full 1024Kb expanded for `ramdrive.sys`. (I do not load any file into the `ramdrive`.) I am now using neither.

I have enclosed copies of my `autoexec.bat`, `config.sys` and `win.ini` files. Please, can you help a newly born 'bluebaby'?

GE Wooding

Well, you certainly seem to be having your fair share of memory problems. Before going into the details of your specific problem, I should probably quickly review the types of PC memory. First of all, there is conventional memory, which is that RAM from the start of the process-or's address space, up to the 640Kb boundary. There can be less than 640Kb conventional memory in a machine, but it is pretty rare. However, 640Kb is the maximum amount of conventional memory that a PC clone can have. Any extra memory which you have then falls into one of the other memory categories.

The simplest form of extra memory (from a hardware point-of-view) is extended memory. This starts at the 1Mb boundary, and extends up from there. Extended memory is only possible on machines with a '286 or better processor. The 8088 and 8086 chips' address space stops at 1Mb, so no extended memory is possible.

Expanded (or EMS) memory, on the other hand, requires more hardware than extended, but as a consolation, can be used with any 8086 family processor. This is because expanded memory resides above the top of conventional RAM, but below the 1Mb addressing limit of the 8088 and 8086 chips. However, while you can have many megabytes of expanded memory (8Mb under LIM 3.2, or 32Mb under LIM 4.0), there is nowhere near this amount of space available. LIM, by the way, stands for Lotus — Intel — Microsoft, the three major players in defining the EMS standard.

Of the 384Kb of space left over in the address map, much of this is taken up with the system and video BIOSes, and the video display memory itself. However, there is virtually always a 64Kb block of free address space, and it is this area of the address space where the EMS memory appears. Since there is clearly much more EMS memory than can be addressed at once through this small part of the memory map, a technique called bank switching is used.

Bank switching divides up the entire chunk of expanded memory into 16Kb pieces, called pages. Since the EMS standard calls for a 64Kb block of memory space, four of these pages are accessible at once. It is up to the EMS driver to switch banks to satisfy memory requests from drivers and application programs. This bank-switching circuitry is the reason for the increased hardware complexity of EMS over extended memory.

MAY 1991

Because of the increased complexity of an EMS memory system, some memory board manufacturers actually implement extended memory, and then use a software driver to make the extended memory appear as expanded — quite valid practice on '386 machines, where the bank-switching hardware is already built into the chip, as part of the memory management unit, (MMU).

'286 machines, on the other hand, have no such bank switching circuitry, and if such circuitry is not provided on the memory board (or motherboard), then all the work needs to be done in software. This incurs a severe speed penalty, because of the amount of switching into and out of protected mode that the processor has to do.

Sorry, I couldn't avoid talking about the processor's modes, so I'd better explain briefly what they are about. The real mode of a '286 processor is a mode which closely emulates the 8086 processor, with a few minor differences. It is in this mode that the processor runs DOS and DOS applications. However, in this mode, the processor can only address 1Mb of memory — the same as an 8086.

To access any extra memory, the processor needs to be switched into protected mode — so-called, because it allows a multi-tasking operating system to protect the various simultaneously-running programs from writing all over each other's memory space. Switching into protected mode is no problem in itself — the '286 has a special instruction especially for the purpose.

However, once we are in protected mode, we run into two problems. The first of these is that the processor's segment registers behave in a different way than they do in real mode. Don't worry too much about the purpose of the segment registers now; suffice to say that software (like DOS), which

is designed to run in real mode won't work in protected mode. Therefore, once the work with extended memory has been completed, we must switch back to real, so that the application program can continue merrily on.

However, once we are in protected mode, we run into two problems.

The second problem is that the processor has no in-built way to switch back to real mode! Fortunately, IBM foresaw this problem, and incorporated a routine in the BIOS, and special hardware on the motherboard, which kicks the processor back into real mode — by resetting it. This is a pretty major sort of operation for any processor, and one which takes a considerable amount of time, (for a computer).

That means that whenever a program needs to access extended memory, be it either pure extended, or an EMS emulation running in extended memory, then a considerable speed penalty is introduced. For this reason, emulated EMS is not an efficient way of implementing this sort of memory on a '286, and extended memory is generally second-best to EMS for the same reason.

So, given the choice of extended or expanded memory on a '286, all other things being equal, expanded is usually the way to go. Windows 2.11 can use expanded memory to increase the amount of space available to run Windows applications, but it cannot use

extended memory. So, your decision to go with expanded seems like a good one.

So, where has your extended memory gone? I see no reason why you cannot have expanded and extended memory in the same system. In fact, I know several people with '286 systems and just that memory configuration. I suspect that when the expanded memory was installed, some DIP switches or jumpers were mis-set.

As a rule, when installing an expanded memory board, you shouldn't have to change anything on the motherboard. If you did change any switch settings during the installation, then try changing them back the way that they were before the memory was added — so that it is set up for 1Mb on the motherboard — 640Kb conventional, and 384Kb extended. This shouldn't effect the way in which the machine operated, but should make the extended memory available again.

If this should fail, then the next step is to remove the expanded memory board, and make sure that the original memory report (640Kb + 384Kb), appears on the screen. Now, double-check the switch settings on the memory board (don't touch the switches on the motherboard), and re-install the card. If the 384Kb of extended memory fails to appear in the boot-up report, then it would appear that your memory card is capable of operating as extended or expanded memory, and this is causing conflict with the extended memory already in place.

If this is the case, then you could try removing the second 512Kb of RAM from the motherboard, and installing them on the memory card. However, this requires that the chips on the card and those on the motherboard are the same size and type. If you do this, set the switches on the motherboard to the setting for

MAY 1991

512Kb, and those on your memory board to allow for the extra half a megabyte which it now has.

Most EMS memory boards have a facility to back-fill main memory, which will allow you to regain the extra 128Kb of main memory, making it back up to 640Kb. Consult your memory board manual for details on how to do this. Your lost 384Kb of RAM will now be available as extra expanded memory, unless your memory board allows you to use some of its memory as EMS, and the rest as extended. This will depend on the exact type of card you have — some are more flexible than others.

As far as software is concerned, Windows 2.11 can use expanded memory if it is available, as can some other applications like Lotus 1-2-3. If, you decide to upgrade to Windows 3.0, then extended memory is the preferable option. This can pose problems if you need extended memory for Windows, and expanded for some other application. In this case, you will need to either switch the memory from extended to expanded and back again when switching applications, or buy more memory, and configure some as extended, and some as expanded.

Coherent Unix

In August last I received a copy of the Coherent operating system from the Mark Williams Company in Illinois, USA. It consists of 4 high-density floppies and a single manual, with 1056 pages. The floppies are not in DOS format, so I could not back them up. It is designed for the AT computer, and requires a hard disk with at least 10Mb of free space.

I had a 40Mb drive, with a 32Mb DOS partition, so I had to re-partition it for a 20Mb DOS partition, and

re-install my DOS files. Then I inserted disk 1 and typed 'begin', and followed the instructions on the screen. It took about 30 minutes to install — the Coherent operating system consists of 1220 files. The manual commands are on disk, and occupy 725 files.

I then tried to back up the four floppies, but had some trouble. With only one drive, I copied the floppy to a directory on the hard disk, but this took 10 minutes! Then it took another 10 minutes to copy the directory to a new floppy. Not like DOS, is it? Finally, I rang MWC and requested four backup floppies, which cost me US\$20.

I have only found one program made for Coherent, called 'rdb', from Robinson Schaffer Wright in Santa Cruz, California, which is on Bourne Shell code. These people say that Coherent is the best operating system on the market, at an unbelievable price. I was going to buy 20/20, but MWC told me that it does not work with Coherent. This spreadsheet needs real Unix, and Coherent has come missing commands when compared with System V.

I am now looking for a Kermit program with a dial option, needed to use my internal modem. I also purchased a copy of c/Database Toolchest, from Mix Software. This includes a 'little database manager' which compiled OK in Turbo C, but Coherent C decided that it had syntax errors on every line.

I would like to hear from anyone using or interested in using Coherent.

*Jon Kitchen
PO Box 355
Mount Hawthorn
WA 6016*

Unfortunately, nobody at our office has any experience with the Coherent operating system. However, if any readers are able to help Jon, or are interested in the system, he can be contacted at the above address. Who knows, we could be witnessing the start of the Australian Coherent Users' Group, if such a group doesn't already exist.

Tech Tip of the Month

THE COMPLETE FAX Portable software uses US format dates — that is MM/DD/YY, which can be a little disconcerting for those who are more used to the Australian convention of DD/MM/YY. Unfortunately, the software installation program does not give this option. A search of the exe files in the package revealed standard DOS 'get date' calls, in files such as _pfmain.exe, _pfm.exe, and _pfout.exe.

However, the most important one is in _pfcover.exe, which automatically inserts the date on the cover sheet of outgoing faxes. Fortunately, it is quite easy to change the date to the Australian format, by following these steps:

First of all, copy _pfcover.exe to another filename, without the .exe extension, for example, pf.x. This is because debug treats exe files differently to other files, and we don't want debug to know that we're working on an exe file.

Then enter debug and load the file to be modified, by typing 'debug pf.x'. Now type 'r cx' to determine the length of the file, (it is less than 64Kb, so BX will be zero).

You now need to locate the 'get date' DOS call, using debug's search facility; type 'S 100 515F B4 2A CD 21'. This will return an address in standard segment:offset form. Now, unassemble the code from that point, using 'U xxxx:yyyy', where xxxx:yyyy is the result of the above search. The commands which you are looking for are 'MOV AL,DH' and 'MOV AL,DL'.

Now type 'E www C2', where www is the address of the MOV AL,DH statement. Then type 'E zzzz C6', where zzzz is the address of the second command, MOV AL,DL. These two changes have now reversed the order of the month and day of the month for all subsequent use in the program.

Type 'W' to write the changes to the file, and then 'O' to return to DOS. Now all that remains to be done is to rename pf.x back to its original name — _pfcover.exe. Of course, all of this should be performed on your working copy of the program, not the original distribution disk. This operation can be repeated for other files, with the exception of _pfout, which only works with the American format.

*ES Webber
Papua New Guinea*

Just to prove his point, Ted's contribution to 'Tech Tips' arrived via his Complete Fax Portable and modified software. And yes, the date was in Australian format. Debug can be a particularly unfriendly beast to use, so be prepared to make a few mistakes if you aren't experienced with it. And remember — don't touch the program on that original disk. Only modify a copy of the program, after you have installed the software.

MAY 1991

June 1991

'386 memory usage

I recently purchased a new '386 (a full-blown job, not an SX), with 4Mb of RAM, a 1.2Mb floppy disk drive, a 40Mb hard drive, and a VGA colour multisync monitor.

My problem is how to best utilise the RAM that is available to me. Most of the programs I have at the moment are quite old, and obviously were not developed in the days of 'extended' and 'expanded' memory. Although I hopefully will be moving up to the 'Wonders of Windows 3.0', and therefore should be able to take advantage of its memory management techniques to best use the 4Mb, at the moment I am only running MS-DOS 4.0, complete with the shell.

Although Your Computer has intensely investigated the topics of, and very well explained the concepts involved with, memory management over the past 12 to 18 months in various articles (which have helped my understanding of the topics very much), not much has been written about practically implementing these techniques on the PC. Could you shed some light on the ways of performing them in MS-DOS, using the respective drivers, as the coverage of this in the MS-DOS manual is absolutely pitiful.

To pinpoint a few exact problems, I have several programs that will not work effectively, and tell me quite often that I don't have enough resident memory to do something or other. This problem occurs especially with a number of the games that I use, for example, BattleChess. Another similar problem occurs when I try to use the 'virtual page' option under the DrGenius software included with the Genius GM-F302 mouse.

I am running DOS 4.0 at maximum DOS functionality, although I did read and take note of your article in the January issue about DOS 4 installation. I have included copies of my config.sys and autoexec.bat files, so that you can gain some insight into the way I have (really badly), applied the memory drivers supplied with DOS already.

Any help that you can give me would

be most gratefully appreciated. Understanding the nature of your column to be such that you discuss common problems that you think a wide group of readers might benefit from, I know that I and many others like me would benefit from your suggestions.

M.E. Krenmayr

First of all, let me assure you that you are not the only person who is not sure how to best use the memory in a powerful machine such as those based in the '386 chip. Of all the components in a PC, the memory sub-system is probably the least well understood. This stems primarily from the number of different addressing modes of this processor (and, to a lesser extent, those of the '286 as well), and the need to switch between these modes to access the memory in different ways. The techniques for addressing these different types of memory have been well documented in Your Computer recently, in the articles to which you refer. However, as you note, these articles have not addressed using this memory with your existing software.

The so-called '640Kb barrier' is a result of decisions taken by IBM in the early stages of the design of its PC machine, especially in that company's choice of processor. In selecting the Intel 8088 chip, the designers restricted themselves to a maximum address space of one megabyte, of which a maximum of 640Kb could be used by system RAM. Although many work-arounds are in common use, especially in conjunction with '286 and '386 processors, the 640Kb limit remains as far as standard DOS programs are concerned, although some of the other pressures on this valuable memory segment can be eased.

On '386 machines, one of the best places to start is to load as many resident device drivers and

TSR (terminate and stay resident), programs into any unused gaps in the memory between the top of conventional RAM (640Kb), and the start of extended memory, (1Mb). However, since there is physically no memory in these areas, some form of memory manager is required to 'pinch' some of the extended memory in the machine, and put it into these unused spaces.

One of the best known memory managers is the Quarterdeck Expanded Memory Manager, or QEMM. Aside from emulating expanded memory on extended memory-equipped '386s, it can also make use of unused spaces in the sub-1Mb memory area, to load drivers and TSR programs. This is achieved through the use of two 'loadhi' utilities which come as part of the package.

Using these utilities, virtually any device driver and TSR can be loaded into a segment of so-called 'high memory', which means that the memory which it would otherwise occupy in conventional (sub-640Kb), memory, is now available for DOS applications. For example, the fastopen, share, nlsfunc, and smartdrv commands in the config.sys file could be 'loaded high', to give you more memory space below 640Kb.

In addition to allowing resident programs to be loaded up above 640Kb, QEMM is also an expanded memory manager — in fact, this is its primary role. Expanded memory is used by some software, most notably Lotus 1-2-3, but on a '386 machine it can be put to good use, by allowing you to run multiple DOS applications simultaneously, using DESQview.

DESQview is a multitasking add-on to DOS, which uses the '386 protected-mode features of this chip (as well as the SX), to allow programs written on the assumption that they have exclusive access to the resources of

JUN 1991

the entire system, to run concurrently with other similar programs. In this mode, each program thinks it is running on an 8086 processor, with its own screen, keyboard, and I/O ports. DESQview then arbitrates between the various processes to determine which one gets keystrokes sent to it, which have access to the screen, I/O ports, and so on.

Many users seem almost apologetic when they say that they are running older versions of software. There is no inherent reason to upgrade software if the version which you have does all that you need it to. If it ain't broke, don't fix it. However, with a '386 machine with this much memory, the software won't make use of the extra memory itself, but there are other ways to use it.

One of these is to use the DESQview/QEMM combination described above, which gives you multitasking capabilities with virtually any DOS application. For example, you could download a file from a bulletin board while sorting your database, and writing a letter with a word processor, all at once.

If you don't think you need multitasking, then another thing which will improve the performance of disk-intensive programs is a disk cache. This is similar to a RAM cache, in that it stores the most frequently-accessed areas of the hard disk in RAM, so that if a sector is needed which has been accessed recently, there is a good chance it is already in the cache, and so no physical disk access is required.

The smartdrive entry in your config.sys file installs such a cache on boot-up, which will definitely improve performance with disk-intensive applications. On the other hand, if your applications are not all that disk-intensive, then the memory could probably be put to better use.

Print caches often make use of extended or expanded memory, but

in my opinion, this is something of a waste. Printers are such slow devices that their caching needs can be quite adequately met by using the hard disk, leaving the RAM free for other applications.

By the way, using Windows exclusively to run ordinary DOS applications, even on a '386, is a waste, in my opinion.

Windows 3.0, on the other hand, uses extended memory, which is the natural way in which any extra memory in '386 machines appears. A memory manager is still required in most cases — Windows uses its own himem.sys driver, although the latest version of QEMM (5.1) is also suitable, and has the added advantage of being able to load drivers up above 640Kb, as discussed above. Some older users of extended memory, like disk caches and RAM disks, work without an extended memory manager, but since there is now no standard way of programs seeing how much extended memory is available, you are usually restricted to using such programs on their own. Current versions of most programs, though, interface to extended memory through an extended memory manager, such as himem.sys.

By the way, using Windows exclusively to run ordinary DOS applications, even on a '386, is a

waste in my opinion. Standard DOS applications gain nothing in terms of functionality when run under Windows, and because Windows runs in graphics mode, it needs to convert text characters which DOS applications write to the screen, into the appropriate graphics pixels, in the correct region of the screen. All of this takes time, which slows down DOS applications when compared with their performance under a text-based multitasker, such as DESQview.

By the way, DESQview is not restricted to just text applications — if any program needs to run in graphics mode, it can do so, and DESQview will emulate text mode for any other visible windows, in the same way that Windows does. However, when any graphics programs are terminated, or their windows are not visible, then DESQview switches automatically back to text mode, for maximum performance.

RAM disks and disk caches are probably the two most common ways to use extended or expanded memory, and of the two, the latter is generally preferable. To use a RAM disk effectively, you have to specifically copy often-used files to the RAM disk, and then copy them back to the hard disk afterwards, if the content of the files has changed. The exception is the temporary files which some programs create. Usually, by setting the TEMP environment variable to point to a RAM disk, some programs which use temporary files can be made to put them in the RAM disk.

A disk cache does all the hard work for you. It monitors disk accesses, and keeps a copy of the most frequently accessed parts of the disk drive in RAM, and it determines whether each disk access can be answered by data already in the cache, or whether a real disk read is necessary. It also takes care of writing any modified

data back to the drive afterwards, when it has time to do it, or when a time-out period has expired. Because the cache takes care of deciding what to store in RAM, you don't even need to know which files are being cached, let alone copy them to a RAM disk and back.

Windows has for some time come with its own disk cache — the smartdrv.sys file alluded to earlier — which is specifically designed to work with Windows. It allows a minimum cache size to be defined, so that if Windows needs more memory, it can recover some from smartdrive for its own use. Smartdrive will also work without Windows, but the dynamic sizing feature cannot be used. The latest version of Norton Utilities (5.0) comes with its own disk cache, and there are stand-alone disk cache utilities available as well. Perhaps we'll look at disk caches some time in the future.

There is one more type of memory present in '286 and '386 machines, and that is the high memory area, or HMA. This is in fact a special case of extended memory — specifically the first 64Kb block. (Actually, it's 16 bytes less than 64Kb, but what's 16 bytes these days?) The beauty of the HMA is that it can be accessed while the processor is in real mode — that mode where these processors emulate the 8086. The HMA results from a quirk of the implementation of the paging scheme in these processors, which we won't go into here.

This means that the HMA can be accessed by these processors as easily as memory below 1Mb, if the software knows how to. In general, the HMA is managed by the same driver which looks after the rest of extended or expanded memory — such as QEMM or himem.sys. Some Windows programs can use the HMA, and some resident programs (such as Program Development Systems' Freeway

network) can use this memory, either exclusively or in conjunction with conventional RAM. Only one program can use the HMA at any one time, however.

When DOS 5 finally appears, it will have the capability of loading a large part of itself into the HMA, leaving about 620Kb of memory free for applications. Hopefully Microsoft won't keep us waiting too long for this — see last month's 'Tech Tips' for a brief look at the features of DOS 5.

Ventura under Windows

I have been a subscriber to Your Computer for the last three years, and have found that its format unmatched by any other magazine. Your magazine has really helped me in the past, not only with the odd school assignment, but also in making the transition from old Apple IIc to IBM '386SX just over a year ago.

When Microsoft brought out Windows version 3 I jumped at the chance to make my computer more 'Mac-like', but I have a dilemma which I have been unable to solve. I cannot get Ventura (version 2) to work under Windows. When I ran the Windows Setup program, it recognised Ventura, and even set up an icon for it. But, when I try to run the program, I get a message saying that I do not have enough memory. My machine has 640Kb of conventional memory, and an extra 384Kb which I am currently running as expanded memory, and I was wondering whether it can be used by Ventura to run under Windows.

J. Meilak

We're glad to hear you like the magazine, and we're always willing to listen to readers' suggestions as to how to improve its content.

Now to your problem. Probably the biggest problem with running non-Windows programs under

Windows is that you end up with less conventional memory than you would running the same application 'barefoot' under DOS. This stands to reason — Windows needs memory for itself, and this must come from somewhere.

Since you're running a machine with 640Kb conventional RAM, and 384Kb expanded, Windows must be running in Real Mode, which does not leave much conventional RAM for Ventura to use. Ventura is a rather big program, and really needs as much memory as you can give it. If you want to use the professional extension features, then you need at least 256Kb of expanded memory. However, if Windows is running in Real Mode, and sees any free expanded memory, it will use it, in an effort to free up as much conventional memory as possible for the DOS application. Unfortunately, this will still leave only about 500Kb of conventional RAM for Ventura, and since Windows has pinched all the expanded memory, Ventura won't get a look-in.

Even if you were to configure the 384Kb as extended RAM, and run Windows in Standard Mode, you won't end up with much free memory. In fact, according to some quick tests which I have just run here, you might even have less conventional RAM than in Real Mode, and no expanded.

My suggestion is to forgo Windows when running Ventura, and run it straight from DOS. With 384Kb of expanded memory, you will also be able to use the Professional Extension, if you have it, with all that it offers. You could upgrade the machine to 2Mb of RAM, and run Windows in Enhanced Mode. This gives about 10Kb less conventional RAM than plain DOS, and it can emulate expanded memory for Ventura. However, Ventura will run slower under Windows than it would under DOS, and I think that anything you

JUN 1991

can do to increase the speed of big programs like Ventura is well worth doing.

There is a Windows version of Ventura 3.0 available for Windows, but it runs slower than the equivalent DOS/GEM version, although it would probably perform better than the DOS version running under Windows, and of course, you can use all the features of Windows which aren't available to non-Windows applications.

Disk recovery

I feel I must comment on the tip on reading the bent disk in March. While your suggestion for resolving this problem may help those without access to the latest software specifically for the problem of recovering data from a damaged diskette, I feel I should tell you about Norton Utilities 5.

Recently a friend of mine received his copy of this program, and decided to put it to the test, with a 3.5-inch diskette which was given to him by someone at work — supposedly unrecoverable — and took the disk home to see if Utilities was as good as the sales people said.

Upon an initial read of the disk, Norton's said that the disk was not formatted. However, when he tried to format it, it told him that there was an unrecoverable physical problem and that he should try something else. Being a person who does not give in, he tried again and again, using different options from within Disk Doctor, and other utilities that Norton 5 has to offer, and after about half an hour, he managed to format this disk, which is now usable.

So, if you are a serious PC user, a copy of the Norton Utilities version 5 is a worthwhile investment. The program has a lot of excellent features for optimising your hard disk, and it would take forever to cover the power of this software.

A. Webster

It would appear that you have misunderstood the purpose of the exercise which I described in the March issue. I was not trying to recover a damaged diskette so that I could re-use it. Indeed, the process which I described required the sacrifice of another disk in addition to the damaged one. The aim was to recover as much data as possible from the damaged disk, without damaging the magnetic coating any further, or damaging the drive in the computer. As it turns out, I was able to read all of the files on the disk without problems, after which I discarded the damaged disk.

Certainly, if some critical part of the disk was unreadable, then I would have had resort to Norton Utilities or something similar, in order to recover as much information as possible and piece it together in a meaningful fashion. Nevertheless, the effort of removing the disk from the damaged sleeve and putting it into a good one would increase the chances of recovering the data — the Norton Utilities (or any other software solution), cannot compensate for read errors induced by a damaged disk surface.

I would put the success of your friend in recovering a floppy disk with a physical defect down to good luck. I'd say that the region in question on the disk was of marginal quality, so that with enough attempts at formatting it, it would eventually format correctly. However, I certainly wouldn't trust any valuable data to such a disk.

Whatever defect was present in the disk's surface coating was present from manufacture, and running any utility software isn't going to remove it, no matter how hard you try. It's a bit like trying to remove some rust from a car body panel by washing and polishing the car every day. Furthermore, if the defect is on an area of the disk which contains the file allocation

table or root directory, then you run the very real risk of losing some important data on the disk — possibly of making the entire disk unreadable.

With floppy disks as cheap as they currently are, I don't waste time with defective ones. If they don't format with no bad sectors, I discard them (once I am sure that it is the disk and not the drive). I guess it depends on the value you place on the data you store on the disk, but I can't conceive any data being worth less than a couple of dollars, and you can certainly get 3.5-inch disks for less than that.

Disk utilities like Norton's certainly do have their place, and are indispensable when it comes time to piece the jig-saw puzzle of a corrupted disk back together. But using the utilities in the way which you describe is just courting disaster. Sooner or later, the disk will lose something, and you'll find yourself once again using Norton's to recover the data from the disk which you 'saved' earlier on.

Hard disk mounting

I have a rather elderly AT, which predates the arrival of small footprint cases, so I am stuck with this big box on my desk, taking up an enormous amount of space. I've considered getting one of those 'feet', and sitting the thing on its side next to the desk, but I'm worried about the orientation of the hard disk. A friend once told me that hard disks were designed for horizontal mounting only, and mounting them vertically would not be a good thing, (I'm not sure whether he meant it would result in physical damage to the drive, or just loss of data). Can you enlighten me?

E. Johnson

Your concern is understandable, but I can assure you that there is

no reason to worry. In the days before the IBM PC, some machines had their hard disks mounted vertically, and others horizontally. IBMs had a horizontally-mounted hard disk, because it suited the orientation of the computer cabinet. All of the early clones had horizontal hard disks, because they just copied IBM.

Nowadays, with the wide variety of enclosures available, especially small footprint ones, vertical hard disks are about as common as horizontal ones. Some enclosures which have horizontally-mounted floppy drives, have the hard drive mounted vertically to better make use of the available space. Those which are mounted vertically are identical to the ones which other manufacturers mount horizontally, so there are no 'vertical' or 'horizontal' drives.

Sure, it is possible that the bearings in a drive might last a little longer in one orientation than the other, but I think that any difference would be small, if it was statistically significant at all. There are many other factors affecting the life span of hard drives, and I'd expect the bearings in most drives to outlast some of the other more fragile parts.

The question of data safety is different, especially with stepper-motor drives. This is because stepper drives don't know exactly where the heads are; they just assume that the heads are positioned where they were told to go. Voice-coil drives have a closed-loop servo system to make sure that the heads are directly over their assigned tracks, so they are always on target.

If you change the orientation of a stepper motor drive, the changed direction of the force of gravity on the heads might move them slightly with respect to the tracks already on the surface of the disk. A voice-coil drive can compensate for any movement because of its

closed-loop feedback head positioning; a stepper motor drive can't. So if your drive is a stepper drive, it is a good idea to low-level format it in its new position, just to make sure that the heads are in line with the tracks. With a voice coil drive, I wouldn't bother, which is fortunate with the proliferation of IDE drives, which in general cannot be low-level formatted without a special utility.

Incidentally, I had an AT with a Miniscribe 3650 stepper motor drive, which I changed from horizontal to vertical mounting about two months after I bought the computer — without doing a low-level format. It never missed a beat.

XT memory expansion

I have a 10MHz turbo XT with 640Kb of RAM and a 60Mb hard disk. One of my main applications is Lotus 1-2-3 version 2, which I have been using for several years. However, as my confidence with the program has grown, so has the size of my spreadsheets, to the point where I frequently run into memory limitations.

Obviously, I would like to upgrade the memory in my machine so that Lotus can handle bigger spreadsheets. However, I am unsure whether I should be looking at extended or expanded memory, or is it time for me to upgrade the entire computer. If the latter is the best choice, what peripherals in my existing system can be transferred to the new machine? Can I just change the motherboard, or would I be better off starting from scratch?

K. Neville

The decision as to which upgrade path you should follow depends on how much memory you want, and how much money you are willing to spend. If you are basically satisfied with the performance of your

system as it stands (except for the memory capacity), then a simple memory upgrade would be the least expensive option.

Upgrading the memory of XT machines beyond 640Kb does not require any difficult decisions to be made — the 8088 processor used in XTs can only address 1Mb of memory in total, which is virtually all accounted for, except for some small areas. Thus, the only expansion option open for such machines is to add expanded (EMS) memory. Extended memory cannot be added, since its address range starts at the 1Mb boundary, which is precisely where the memory addressing capability of the 8088 ends.

Expanded memory, on the other hand, only occupies 64Kb of address space, and a suitable space exists between the top of conventional RAM (at 640Kb), and the top of addressable memory, at 1Mb. So expanded memory is the only choice for XT machines, and it is also the type of memory that Lotus (and most other memory-hungry packages), requires.

When purchasing a memory card, make sure that it is capable of being used in an 8-bit machine, which your XT is. Many memory cards are 16-bit cards, but can be switched into 8-bit mode, either manually or automatically. If you get such a card, and upgrade at a later time to a 16 bit machine, then you will still be able to use the memory board should you want to.

If, on the other hand, you want to take the opportunity to increase the speed of the machine at the same time, then you have three other courses open to you. The first of these is to install an accelerator board, containing a more powerful processor. These boards are available with '286, '386, and '386SX processors, and usually contain an on-board cache of some sort. However, unless the board has its own memory, you won't be able

JUN 1991

to expand beyond the 640Kb limit which you have already reached, without using EMS memory. This is because the XT expansion bus doesn't have the address lines needed to address more than 1Mb of memory. Also, your system will still be throttled by the 8-bit bus in the system, so it will never reach the performance of a system designed around the particular chip from the ground-up.

A motherboard transplant is the next option. Again, depending

on the physical constraints of your box, you have the choice of a '286, '386 or '386sx processor. You can get these to fit most XT boxes, and they allow you to have a 16- or 32-bit data path to the memory.

However, this doesn't address the issue of the hard disk's data transfer rate. Since the hard disk controller in your XT is only an 8-bit card, it won't be able to transfer data as fast as a proper AT controller, which uses that machine's 16-bit bus. So you can

add a new hard disk controller.

So, on balance I'd have to say that if you want a faster machine, start from scratch (except, perhaps for the hard drive itself, if it is reasonably fast). It will ultimately work out not that much more expensive than upgrading your existing system, if at all. On the other hand, if you are happy with the speed of the system as it is, and don't plan on buying any graphics-based software, such as Windows, then just get the memory board.

Tech Tip of the Month — Indexing images

I RECENTLY BOUGHT a new mouse, and a copy of Microsoft Paintbrush. After playing around with the product for a while, I soon found that I could not remember which image was in what PCX file. I experimented with the new version of Isys, and came up with a technique that may be of interest to other readers. My particular problem was very simple, but the technique could be used for all sorts of applications.

Since this technique links text and graphics, I think it is fair to refer to it as 'multimedia', although it is a very different beast to IBM's multimedia software that was recently reviewed. I understand that a relatively small percentage of your readers would be users of Microsoft Paintbrush and Isys, but I hope that you will consider this 'tip' of sufficient interest for publication.

As stated above, there is much interest in multimedia these days, with many articles appearing in the press about packages that combine graphics, sound, and text. However, there are many application areas that do not need the level of sophistication that these provide, simply requiring the ability to link a graphical image with a text description. This could be used in a variety of ways, such as a photograph catalog for graphic artists, linking scanned articles with abstracts in a technical library, or displaying a picture or drawing of a property to go with a real estate agent's description.

In these examples, the description is used to retrieve the information — such as all photographs featuring sunsets, with

tree or bird silhouettes. The descriptions are held as normal word processing documents or ASCII files, and located using Isys, which is a full text-retrieval product.

Isys version 2.0 has an 'activate document' feature, which may be selected when viewing a retrieved document, from the 'document handling' menu. Normally, this is set up to invoke the appropriate word processor, and load the document that is on the screen. However, the feature can be tailored to perform other functions, such as displaying an associated graphics image. For example, a series of PCX files may be created and maintained using Microsoft Paintbrush, and stored in the pbrush directory. Each image has a corresponding description file, and the link is in the naming of the files. The description of sun-set01.pcx is in the file sunset01.txt, and so forth. The text files are indexed as a normal Isys database, and have a one-line descriptive title at the start of each document. The DDA (dynamic document activation), definition has to be changed, to activate Microsoft Paintbrush, and load the PCX file, instead of calling the text editor and loading the TXT file. The following comprise the DDA definition:

```
Directory \PBRUSH
DOS Command PAINT
Parameters
Keystrokes {PAUSE10}{NAME}.PCX{ENTER}
```

When activated while browsing a document, Isys will move to the pbrush directory, invoke the paint program, type in

the name of the PCX file, and press the enter key. The }PAUSE10{ keyword causes a one second pause while Paintbrush loads, so that the keystrokes are not lost.

In use, the searcher enters a query such as 'computer'. The document list shows the titles of all the documents containing the word 'computer', such as 'PC with exploding screen', 'woman carrying portable', 'laptop computer next to briefcase', and so on. To read a document, the user highlights the title, and presses the enter key. To view the image, the user presses F2, for the document handling menu, then 'A' to activate. At the end, the document list is displayed again.

Microsoft Paintbrush uses a lot of memory, which limits the complexity of the images that can be displayed, since Isys is also resident in memory. One option is to use a public domain PCX file viewer, such as Picem, to display the images. In this case, the DDA entries should be:

```
Directory \PBRUSH
DOS Command PICEM
Parameters {NAME}.PCX
Keystrokes
```

This technique can be used with a number of graphics packages, for a variety of applications. The format of the image will depend upon the software used to create it. However, there are public-domain viewers, if the original software is too memory-hungry to be invoked from Isys.

Tom Keech

July 1991

Missing utility?

I have been using DOS 3.2 at work for quite a number of years, and have re-programmed some of the Function keys on that machine. This DOS version (not MS-DOS), utilises the `ansi.sys` driver with a file called `esc.com` to program the Function keys using Escape sequences.

Recently I bought a computer of my own, and had my first look at 'real' MS-DOS, version 4.01. I tried to re-program some of the Function keys using Escape sequences, but was only rewarded with the 'bad command or filename' message. I then noticed that there was no `esc.com` file, so I checked the original floppy disks, but the file wasn't on any of them.

I decided that either MS-DOS doesn't use an external `esc.com` file, or that I have an abnormal version of MS-DOS, considering that the DOS version I use at work does have `esc.com`. I then copied `esc.com` from my work machine onto the machine with DOS 4.01, and the Escape sequences worked fine — even F11 and F12 could be re-programmed.

So my question is: why doesn't MS-DOS 4.01 have an `esc.com` file, or if it does, where's mine? Also, how do you reprogram the Function keys in MS-DOS 4.01 without an `esc.com` file? By the way, the DOS 4.01 manual specifically states that Escape sequences are used to reprogram the Function keys.

Ron McGregor

The file `esc.com` is not a standard part of DOS, but is an extra utility supplied by the manufacturer of the system. Often system makers obtain an OEM license for DOS from Microsoft, so that they can customise it for their own machines' specific requirements. They also often add their own utilities which they think may be particularly useful to users of the machine. `Esc.com` would appear to be such a utility.

If you don't have a utility like `esc.com`, then the Function keys

can be re-programmed by echoing Escape sequences to the screen, with `ansi.sys` installed. This was covered in detail in past installments of 'Tech Tips' — in the August 1990, December 1990, and April 1991 issues.

More on the great memory mystery

I'm not sure whether G.E. Wooding is any the wiser having read your reply to his letter in the May 1990 issue. Even so, others may be suffering a similar fate, without any consolation from the item.

I'd say the most likely cause for the sudden loss of 384Kb of RAM is a feature known as shadow RAM. Many motherboards automatically reserve 384Kb of RAM for BIOS shadowing whenever more than 1Mb of RAM is installed. Whilst the shadow feature can be enabled or disabled under CMOS setup, the memory is still reserved, even if unused.

Some motherboards allow any unused RAM in this area to be returned to the general RAM pool, but many do not. Some motherboards will only count 1.664Mb during boot-up, while others will show the full 2.048Mb. Either way, DOS' `mem` command will only show a total of 1.664Mb, made up of 640Kb base RAM, plus whatever mix of extended and expanded memory you use.

If the motherboard allows you to allocate some RAM for EMS (expanded) memory — those with the Chips and Technologies NEAT chipset do — and you do not load the `ems.sys` device driver in `config.sys`, that RAM will also be missing from the MEM display. Usually in this case, the system will complete the `mem` command, display the RAM map, and then the system locks up.

Shadow RAM is merely the system and video BIOS ROMs copied into (faster) RAM. This generally speeds things up a bit — up to 5 per cent on some tasks. However, I have found that Windows 3.0 to be a bit fussy about video

BIOS shadowing, so it's best not to use it unless you carefully evaluate it, and are prepared for a few system hang-ups in the process.

Assuming the extra RAM Mr Wooding installed was on the motherboard, and the motherboard has a hardware EMS setting, the extra left-over RAM can indeed be split between EMS and extended memory. NEAT chipsets usually come with Quickset software, which makes this task simple. Usually EMS is allocated in 512Kb chunks, leaving half a megabyte for extended.

If Quickset is not supplied, the delicate task of changing settings using the CMOS setup routine is required. Read the screen prompts carefully and avoid changing anything other than EMS enable/disable, and the EMS size. On most systems using the C&T chipset, an AMI BIOS is used, allowing reversion to a default CMOS setting by holding down the Insert key whilst powering up.

Now things get really tricky. Windows 2.11 is a real dog to set up properly, and Windows 3.0 is, on average, much easier. Windows 2.11 only uses expanded memory, but the usual implementation of a Chips and Technologies chipset is 'small frame' EMS, which doesn't help much. If large frame EMS — that where part of the 640Kb of base memory is used up by EMS — is available, use that. This allows fully-conforming LIM 4.0 software to load executable file data into EMS. Windows 3.0 also works best with EMS if it is large-frame.

If large frame is not available, I would suggest switching to Windows 3.0 and using extended memory; it's a lot quicker and more versatile. Here, SmartDrive supplied with Windows 3.0 is the best way to use the extra RAM.

If some applications need expanded memory, and you set the system up with extended memory to run Windows 3.0, try using a product called Above Disk. This is an expanded memory emulator which can be switched in or out via a batch file command. The only snag is that it seems to leave a bit of itself behind on each switch-out. Over time this adds up, and reduces available base RAM.

JULY 1991

One last point on expanded and extended memory with Windows 2.11 and 3.0. I have found that Word for Windows 1.0 running under either the Windows 2.11 run-time supplied, or Windows 3.0 using extended memory, is not very robust in its memory usage. Documents over about 100 pages in length seem to cause it some headaches, and frequently sulks, requiring a warm boot to escape. However, WFW 1.1 under Windows 3.0 on a 1Mb '286 can manage the same 100 page document copied onto itself many times, to produce over 1600 pages without a hitch. It's a bit slow though!

Lawrence Glynn
Sage Advice

Thanks for putting down on paper some advice which I'm sure many users of '286-based machines will find extremely useful. One of the problems with running a column like 'Tech Tips' is that many letters we receive asking for help do not include all the information required to accurately diagnose the problem. This oversight is usually the result of the writer of the letter assuming some knowledge about their particular system, or omitting details which they feel are not relevant to the problem — after all, a letter describing every possible detail of the system would clearly be rather long, (although Word for Windows would appear to be a good word processor to use to create it).

What wasn't specified in the original letter in question was whether the extra RAM added was on the motherboard, or on an extra EMS memory card. I assumed the latter, but in hindsight, your assumption of the memory being mounted on the motherboard seems more reasonable.

Hopefully, with my original reply, and your reply above, we've pretty much got it covered as far as '286 machines are concerned.

Booting DOS and Xenix

I note in the April issue that you describe a procedure for changing the active partition table, to enable DOS or Xenix to be booted on the same machine. While the procedure probably works, there is a far simpler and safer approach.

The machine should have DOS installed first, and then Xenix, (though this is not strictly necessary). In any case, the Xenix partition should be made the active (bootable), partition. This should then not be changed. By default, the machine will now boot Xenix on power-up or reset. However, as it does, it puts a 'boot:' prompt on the screen, to allow specification of the copy of Xenix to be run.

If you type nothing, this will time-out after 60 seconds, and boot the standard Xenix. However, you may also enter 'dos' at this prompt, whereupon it will boot from the DOS partition (without altering the partition table). Thus, there is no need to fiddle around with floppies, fdisk, or any other utilities.

Clyde Smith-Stubbs
Hi-Tech Software

It's odd that our original correspondent didn't notice this feature on his particular implementation of Xenix. All I can surmise is that he was running a different version, which doesn't have this feature, or it is not mentioned in the documentation. My Unix/Xenix experience is entirely from non-PC machines, so I wasn't aware of it either.

Copy-protected software

I have an IBM AT-compatible with 5.25-inch 1.2Mb and 360Kb drives, and a 3.5-inch 1.44Mb drive, and a hard disk. I'm looking for a copy program so that you could select the required format to

copy from or to, and which would overcome the copy protection on the original disk. Also, could you advise of a program that would bypass key-disk protection as well.

While I agree with the manufacturers putting such protection on their disks, I have young children using this computer and my practice has been to copy the original disks, and let the children use the copies. If for some reason the copy is damaged, there is no loss.

On some games you buy you are given the option of purchasing a backup, but for some of us, the cost of the original is stretching the budget. I have made some inquiries about CopyIIPC and Copywrite, but different distributors have given me contradictory information and I am rather confused.

E. Bond

Copy protection of disks is a nasty habit which we had all hoped software publishers would have grown out of by now. While very little business-related software is protected nowadays, sadly the same cannot be said for games. Most business software either has no protection (other than legal protection offered by copyright law), or uses something much less offensive than just making the disks uncopyable.

The writers of games software seem to be of the opinion that there is no legitimate reason to copy their disks, and any such copying would be a breach of their copyright. The problem is that the medium which they distribute the software on is far from infallible — the heads of floppy drives rub on the disk's surface whenever it is spinning, so it is undisputable that the disk will eventually wear out. Perhaps they hope that, by that time, you will have tired of the game, and not bother with it.

This arrogant attitude can only be dealt with by putting pressure on the manufacturers of games, to drop their protection, or use some

other form of protection which won't render your investment worthless if a disk fails. This was why publishers of business software eventually dropped disk-copy protection on their products, added to the fact that there were already solutions in the public domain for virtually any protection scheme that they tried.

Copy protection of disks is a nasty habit which we had all hoped software publishers would have grown out of by now.

The two programs which you mention are both well established, and can copy the majority of protected disks. It is impossible to write a copy program to counter every possible copy-protection scheme, due to hardware limitations in the standard drive controllers. However, any reputable dealer should allow you to take one or two of your protected disks into their store and try the program out, before shelling out your money.

I haven't tried Copy II PC, but an early version of CopyWrite which I used some time ago could handle everything I threw at it, with one exception. This was a specialised engineering program which has one of the best (worst) copy-protection schemes I've ever seen. Even the Copy II PC deluxe option board,

which uses special hardware to simulate that in dedicated disk copiers, can't handle it — although I'm not one to give up easily. That's not to say I've tried every game around — I'm a confirmed games non-addict — but I'd say that your chances are pretty good that it will copy your games. The publishers of this program also update it from time to time, to cope with new protection schemes that appear. It's not unlike the battle between the scum who create viruses, and those who write anti-virus software; I think of disk copy-protection as a sort of virus.

Of course, if enough people sent nasty letters to those responsible for this scourge, then perhaps the problem would go away once and for all. Now the high-profile Business Software Association has successfully prosecuted a few particularly bad offenders, people are much less willing to flout copy-right law, and make illegal copies of software. This can only lessen the need for copy-protection on disks.

RAM disk follow up

I read with interest your suggestion to B. Collins in 'Tech Tips' in the March issue, that an AT EMS board should be left behind when upgrading, because I was in a similar situation recently.

I upgraded my XT (which had an 8-bit EMS board) to a 16MHz '386SX. Since the EMS board was only new and I clearly would not recoup my costs if included in the sale of the old machine, I decided to install it in a spare 8-bit bus connector in the new machine. A little fiddling was necessary to solve a memory conflict between the 512Kb VGA card and the driver for the EMS card, but eventually the whole lot worked perfectly.

I now find that, with my system anyway, Windows runs a whole lot faster with the 2Mb of EMS set up as a disk

cache (using smartdrv.sys, supplied with Windows) than without. Word for Windows, for example, takes about two thirds the normal time to load and speed-wise it is now limited by the CPU speed, not the hard disk. I am limited to Standard Mode in Windows if I include the EMS driver in config.sys, but that is easily remedied if I need to run in Enhanced mode.

It may also be worth mentioning that MS-DOS includes an EMS driver for '386 machines (EMM386.SYS), for those who need EMS outside Windows, but can't afford QEMM and the like, (even though they are clearly superior). Note this driver is included with DOS 4.01, not Windows.

In conclusion, I would suggest that installation of the EMS board is worthwhile even if it is only used for a disk cache — even 8-bit transfers are probably faster than most hard disks.

Linda McGarry

Indeed they are faster. In fact, I've yet to see a hard disk which is faster than even the slowest memory board. Your suggestion is probably the best use for an old EMS board when all your software needs extended memory — such as Windows does. However, the problem which you describe limiting Windows to Standard Mode could be a problem for some people — I'd personally be lost if I couldn't fire up Norton Commander in its own window.

However, using the otherwise useless (as far as Windows is concerned), EMS memory for a disk cache frees up more of your extended memory for Windows itself. While it will be slower than a disk cache in 16-bit extended memory, it will definitely be faster than any non-cached hard disk.

I guess the thing to bear in mind is that virtually each situation is different — both in terms of equipment configuration, and software requirements — and what may be the best solution in one case may not be so in another.

JULY 1991

Mac disks in a PC

I have just finished reading John Hepworth's article on PC tools in the April 1991 issue of *Your Computer*, and in it he mentions a piece of hardware, a Central Point Copy II Deluxe Option Board, that may solve a problem for a friend of mine.

My friend works part time in an office with an Apple Macintosh running Microsoft Word. At home I have an IBM compatible AT, also running Microsoft Word, version 5.5. Owing to current commitments (children), my friend must be at home more than the boss would like.

How do we transfer Word files between the IBM and the Macintosh? Will the Central Point board do the job, and, if so, then what exactly do I need to help her out? What other software or hardware is required?

Ian Phillips

The Copy II PC Option board is probably the solution to your problem. It is, in fact, an add-on disk controller, which allows your existing floppy disk drives to do things which the standard PC controller can't do. One use for this is to copy protect PC-format disks, which often use special tricks that the standard PC controller can't replicate.

The other use, which will be of more interest to you, is its ability to read and write Mac-format disks, when used in conjunction with a suitable 3.5-inch disk drive. The manufacturers of the board recommend virtually any 720Kb drive, or certain makes of 1.44Mb drive — their favourite is Teac.

This card mounts in a free card slot in an IBM PC, XT or AT compatible, and connects between the existing floppy disk controller and the drive(s), using a supplied cable. To read and write Mac disks, at least one of these drives must be a 3.5-inch drive, as described above. The floppy drives continue

to operate normally under DOS, but when the special software supplied with the card is invoked, the card takes over control of the drives from the standard controller.

The software takes care of reading and writing Mac-format files, folders (directories), and the like, and supports both the 400Kb and 800Kb disk formats of the Mac. Not supported is the new high-density 1.4Mb Mac format, but this is not a problem, as I will discuss below. The software does not take care of file format translation, but if you are using packages on the Mac and PC which have compatible file formats, such as MS Word, then there is no need to do any conversions anyway.

If the Mac which your friend is using has a high-density drive (Apple call it an FDHD drive), then it can read and write both 720Kb and 1.44Mb MS-DOS format disks, and so no extra hardware is necessary in the PC. A utility on the Mac called 'Apple File Exchange' takes care of working with DOS disks, or you can get a copy of Dayna Communications' 'DOS Mounter', which allows DOS disks to appear on the Mac desktop as would a normal Mac disk.

The Copy II PC Option Board is available from Software Express, on (03) 663 6580. DOS Mounter is distributed in Australia by NetComm, who can be contacted on (02) 888 5533.

Laptop scanner

I have in my possession an Artec colour scanner — model number A256C — which I would like to use in conjunction with a Toshiba T1600 laptop computer, if possible.

At this point in time I am not greatly concerned with being able to use the full colour spectrum of the scanner

— grey tonings would be sufficient if colour was not possible.

Peter Wing Tang

While I am not familiar with the particular scanner you are using, I presume that in common with most scanners, it has some sort of interface card which plugs into the expansion bus of the PC.

Herein lies the root of your problem, and there are two solutions which I can think of.

The first is to buy Toshiba's desktop docking adapter, which provides (from memory) two standard expansion slots, one of which could be used for your scanner card.

However, this is not an inexpensive option, as the docking adapter provides much more than just a couple of expansion slots.

The use of a single-card adapter, called the Wonunder, would be a somewhat cheaper solution. This provides room to mount one three-quarter length expansion card, and mounts underneath the laptop's enclosure. The Wonunder can be obtained from XLtech on (02) 975 2111.

As far as colour scanning is concerned, I can see no reason why this should not be possible with your machine, although, of course, you won't be able to see any colour results on the internal screen.

This will depend on the software which accompanies the scanner, but since the display adapter in the T1600 emulates a colour adapter (as do all Toshiba laptops), then as far as the software is concerned, the screen is colour, although, of course, you see the colours as shades of grey.

To see the colour rendition, you will need to connect an external monitor to the appropriate port on the computer.

3.0 problems

Your magazine has been of tremendous help to me, and I hope that you will be able to resolve three problems I have, relating to Windows 3.0. I am running a Compaq 386N, with 2Mb of RAM.

The first problem is that when I turn my computer on, a message flashes on the screen declaring 'out of environment space'. What does this mean, and how do I overcome it?

The next problem is that I installed Lotus 1-2-3 V3.1 to run under Windows, but when I double-click on the Lotus icon I get a dialog box saying: 'This application has insufficient memory for its display. Check PIF settings.' The Lotus icon appears at the bottom of the screen, and when this is double-clicked, Lotus starts with no apparent side-effects. I have changed some of the memory settings in the PIF file, but to no avail.

Are these two problems related, and if not, how do I solve the second problem?

My final problem is related to my Microsoft mouse. It works properly when working with Windows applications, but when I try to use non-Windows applications, such as 1-2-3 and Xtree Gold, the mouse doesn't work. When I exit the non-Windows application and go back to Windows, the mouse works once again.

Robert Summers

First things, first. The environment space is an area of memory reserved by DOS for storing variables, such as the command path, and pointers to various directories used by programs. By setting an environment variable when the system is booted up, or from a batch file, a program can find the directory which contains other files which it requires, without this having to be hard-coded into the program's main executable file.

These variables each (naturally) take up space, and if there are too many of them, or they are too long (say, you have a lot of directories in your path), then the special reserved area of memory may not be big

enough to store all of them — hence the error message. The default value of 160 bytes is not very big, but you can increase this up to a maximum of 32Kb if required.

The size of the environment space is specified when the command interpreter is loaded, in the config.sys file. The statement:

```
shell=c:\dos\command.com
c:\dos/e:nnnn /p
```

tells DOS to load command.com in the dos directory as the command interpreter, with an environment space of nnnn bytes, and make it permanent. If you don't have a 'shell=' line in config.sys, and DOS is in a directory called 'dos', just add the above line to your config.sys file somewhere, replacing nnnn with your desired environment size. If there is already a 'shell=' line in config.sys, just add the /e:nnnn parameter to it.

The required environment size is best found by trial-and-error. Find the lowest number which doesn't produce an error, and add a few extra bytes to it, to allow for extra variables to be defined later on, if desired. I use an environment size of 256 bytes, which is adequate for my needs. Don't set it too large, as it takes memory away from the pool of memory available to run programs in.

The next problem is related to the way in which Windows manages memory for multiple applications, and is not related to the previous problem. In order to protect non-Windows applications from writing directly to the screen, and messing up the display of other programs, Windows re-directs all screen accesses to a screen buffer which it maintains for each non-Windows application running on the system.

Because different screen modes (such as text, mono graphics, and colour graphics) use different amounts of memory, Windows needs some means to figure out how much memory to allocate to the screen buffer for each application. Windows

could just allocate enough memory for the most memory-hungry video mode, very wasteful if the application doesn't use that much.

The solution is to store information in the PIF file for that program, describing the amount of video RAM needed by that application. This is accessed through the PIF editor (by clicking the 'Advanced' button if running in 386 Enhanced Mode). Select the option which represents the type of display which the program (Lotus), expects to see — in your case, 'High Graphics' would be the correct choice. Also select 'Retain Video Memory' so that when the program is not running in graphics mode, other programs can't pinch its graphics memory.

Your mouse problem is the result of a dilemma that the writers of Windows (and other multi-tasking systems like DESQview), face — whether to give mouse information to the operating environment, or to the applications themselves. In the case of DESQview, the writers decided in favour of letting each application use the mouse if it needs to, otherwise DESQview can use it for its own purposes. This course makes sense for applications which do their own input and output, as most standard DOS programs do.

However, Windows was designed to take care of, amongst other things, all input and output for applications, including the mouse. This information is then passed to the relevant application by Windows. However, non-Windows applications do not know how to obtain mouse information from Windows, and so they cannot 'see' the mouse when running in a window. The only way to allow a non-Windows application to use the mouse is to run it in full-screen mode, by selecting that option using the PIF editor.

In this mode, Windows can be sure that the mouse is not over some other window which the foreground application shouldn't touch, and so can safely turn over control of the mouse to that application.

JULY 1991

August 1991

Floppies and mice

I would be grateful if you could resolve or explain a conflict of information about 5.25-inch floppy disk drives. I have moved from an XT with a 360Kb floppy drive to a '386SX clone with a 1.2Mb drive, and I am still often writing on this drive to 360Kb formatted disks, to read on the 360Kb drive in the XT. I keep seeing suggestions in your pages that this can cause trouble, but my friendly, helpful dealer, has told me that this is nothing to worry about, and to date, I haven't had any trouble. Can you tell us whether this is a real problem, and what about 3.5-inch drives? I use Norton's Safe Format to format the disks.

I also have a mouse called an Agiler Mouse, that came packaged with my '386SX. It claims to run in both Microsoft and Mouse Systems modes, whatever that means. It works quite satisfactorily with Windows 3, but only if I hold down its left button during the entire power-up procedure, including the RAM and system check. This works, but takes time, and it's getting to be a pain, so do you know of any way around it?

PC Owen

Your question about compatibility between double- and high-density drives is a common one, and has been causing many people headaches since IBM introduced the AT, with its high density drive. Double density disks contain 40 tracks on each side of the disk, at a density of 48 tracks per inch. Thus, all of the tracks occupy slightly less than one inch of the disk's radius.

High density disks achieve an increase in capacity for two reasons. Firstly, the number of sectors of data on each track is 15 rather than 9. This is no problem, but it doesn't account for all of the increased capacity of the higher density disks. The other factor in the equation is that the number of tracks has been doubled, to 80.

However, these 80 tracks occupy the same space on the disk as the 40 tracks of a double density disk — with a density of 96 TPI. In order to read and write these narrower tracks without messing up the adjacent ones as well, the heads in the drive have been made narrower.

A high density drive can read and write 360Kb disks, by stepping the head to alternate tracks, so that it effectively becomes a 40 track drive. Unfortunately, the head remains the same size, so any tracks written on this drive will be the narrow high density tracks, with large gaps between them.

Now, if the drive writes such tracks over the top of the wider tracks, not all of the previous track will be overwritten by the new track, and when a double-density drive reads this track, the signal it receives will be a mixture of the old and new ones — the most likely outcome will be corrupted data.

So, why does the problem arise only sometimes? For starters, if the disk is new before being used in the high density drive, then there is no signal already on the disk to interfere with the desired one. In addition, certain brands of high density drives do not appear to experience trouble — Teac is one such drive. Presumably this is due to higher recorded signal strengths or slightly wider heads. Whatever the reason, 360Kb disks written by Teac high density drives have never given us any trouble.

I can't claim to have experience with the particular mouse you describe, but the power-up sequence sounds similar to one I've seen for the Mighty Cat, (funny name for a mouse). If you leave it alone, it powers up in Mouse Systems mode, or if you hold down the button, it emulates a Microsoft mouse.

If your mouse is the same as this one, run the Windows Setup program, and change the mouse

type to Mouse Systems. Now, you shouldn't have to hold down the button when re-booting. The Microsoft and Mouse Systems standards are the most common for mice, and most emulate one or the other, or in your case, both.

Windows and disk formats

I have two questions for you — one will be a piece of cake, and the other, I'm not sure about. Before getting to them, I'd like to say how interesting I find the 'Tech Tips' column — it's a learning curve all the way! I hope you will carry it on.

I, and thousands of others, I expect, have an XT clone (an Amstrad PC1640, with EGA display and hard card), which has, and is giving, faultless service. I'd like to be able to use Windows 3.0, and from what I read, an AT-class with a '286 processor appears to be a requirement. Can the XT successfully and reasonably be upgraded and, if so, would you publish a 'how to' article?

Secondly, could you please explain the differences between high and low level formats, and what each does? I thought a format was just that — wipe everything off the disk and start again!

Thanks for an extremely informative magazine; I look forward to it every month.

Rod Dent

Thank you for your interest in the column — it is always gratifying to get bouquets. It's only letters like yours which keep us going sometimes, and we'll certainly keep the column going for as long as it serves a purpose.

While you can, in theory, run Windows (in Real Mode) on an XT, the performance leaves much to be desired. Also, you won't have much memory to run any actual applications after loading Windows itself. A processor upgrade of some

AUG 1991

sort is definitely a good idea, but you would probably be better off getting what you can for your present system, and upgrading the entire system.

Simply upgrading the processor will only serve to highlight other bottlenecks in the system, such as the hard disk, and memory capacity. A new motherboard would upgrade both the processor and memory in one hit, but physical constraints may limit the number of suitable choices. The power supply may not be up to the increased load imposed by the needed extra memory.

A '386SX processor would be preferable to a '286 too, because although there is not a lot of '386-specific software around at the moment, more is coming out all the time, and that chip's far superior memory management is advantageous for running Windows. Consider 2Mb of memory a bare minimum for Windows.

All things considered, it is hard to recommend upgrading the motherboard of your existing system, since so many other system components will need upgrading as well if the full potential of the new motherboard is to be realised. Seriously consider trading in the machine you now have on a '386SX, and it will serve you well for a number of years.

Formatting a hard disk, as you know, is a two-step process. The first of these, the low-level format, takes place on a virgin disk, and defines the tracks and sectors which will eventually be used to store the data. As the tracks are laid down, the magnetic coating of the disk is tested for data retention. Any sectors which do not pass this test are marked as bad, so they will not be used to store data, which could be lost at a later date.

Most hard drives are already low-level formatted from the factory, and in the case of IDE drives, you shouldn't try to perform

a low-level format unless you have software which is specifically designed for these drives.

Low-level formatting cannot be performed by any standard DOS program, but many systems come with such a utility either on a diskette, or with it built into the BIOS ROMs.

A high-level, or logical format, is just setting up the file allocation tables and root directory, so that DOS can make sensible use of the space on the disk. This is performed by the DOS format utility, although when this program is formatting floppies, it is actually doing a low-level format as well.

Speed up your IDE drive

I've got an IDE drive which is a bit of a bottleneck in my otherwise satisfactory '386SX system. How difficult would it be to add a hardware cache? And, could you please explain what IDE signifies?

Jon Killen

As most of you will already have noticed, the intelligent drive electronics (IDE), interface has become the standard of choice for all but the highest performance machines. The incorporation of the entire disk controller into a few VLSI chips actually on the hard drive itself, has eliminated the need for an external controller.

The trouble is, you're limited to the controller that the drive manufacturer thought you should have — there is no opportunity to, mix and match drives and controllers. Many larger capacity drives have a cache on-board, but this is usually only 32- or 64Kb, enough to cache one track to allow a 1:1 interleave. Caching controllers are a popular way to improve the performance of drives, and unlike

software disk caches, they do not impose any additional workload on the main processor. The trouble is, with an IDE, you can't just unplug the controller and plug in a caching replacement.

I've just been playing with HiCom's answer to the problem — the DriveCache, a 4Mb caching IDE controller. It doesn't actually replace the controller on the drive — that's not really feasible, but works in conjunction with it. It replaces the paddle card in machines which don't have the IDE interface on the motherboard, and provides the floppy disk controller functions as well, so you can just remove one card and plug the new one in.

If you have one of those combo cards, with serial and parallel I/O ports on the card as well, then you'll have to disable the hard and floppy disk controllers on that card if you want to leave it in place to perform I/O tasks. Either that, or pull the card out completely and get a separate I/O card.

If the IDE connector is on the motherboard, and cannot easily be disabled, then the caching controller can be connected to this through an optional jumper cable. On-board memory can be anywhere between 512Kb and 4Mb, and uses standard SIMMs.

I tried the controller out in the Syncomp Mega 386i that is serving as my main 'test' machine at the moment, and encountered my first problem almost immediately — the board has two small tantalum capacitors mounted very close to the edge closest to the front of the machine, which crashed into the card guides there. I'm sure this problem is not unique to the Syncomp, and the capacitors should be moved on any future revision of the board — there's no reason for them to be this close to the edge of the board.

Once installed, the computer booted up as per usual, except for a

AUG 1991

September 1991

same size. If you have a 4Mb machine, with Smartdrive doing the disk caching, then you'll notice a dramatic improvement in speed with the DriveCache, and all of the 4Mb available to Windows.

Getting the disk cache out of system memory and onto its own card allows Windows more program workspace, which means less disk swapping, and any disk activity which still needs to be performed will benefit from the speed improvement offered by the external cache.

Of course, buying a few more megabytes of RAM will give similar results, but performance will probably not be as good as the DriveCache. Also, many machines come with the motherboard fully populated, so to install any extra memory you'll need to shell out extra for an expansion card, which could end up costing you as much as the DriveCache.

If you have a '286 processor with an extended-memory cache, then the performance increase will be even more dramatic, since the processor won't need to be switching in and out of protected mode all the time.

Of course, if you opt for the DriveCache instead of upgrading your system RAM, then you can't use the memory for anything except as a disk cache. On the other hand, you don't have to worry about software incompatibilities — because DriveCache is hardware not software, it doesn't have to co-exist with other software — it will work just as well, I'd imagine, with OS/2 or Xenix, as it does with DOS.

In the short time I've had the cache, I have gotten used to the extra speed of the computer, and I can't bear the thought that I'll have to send it back. The sample came from Hi-Com Uni-tronics which is fitting the cache to its high-end machines — and is priced at around \$1600 with 4Mb of RAM, (02) 525 8318 if you're interested.

QRAM and extended memory

I have a memory problem. That is, my computers have memory problems, (I do too, but that is another story). I have an NEC Powermate '286 in the office and a Samsung SPC 6100 ('286) at home. Both machines have 1Mb of memory, with 384Kb available as extended memory.

I recently tried to use Quarterdeck's QRAM to make use of this 384Kb, but, on boot-up, I kept getting a message that said QRAM had nothing useful to do. The QRAM manual says that it can use extended memory if it complies with the XMS specification.

If QRAM can find nothing useful to do, does that mean I have done something wrong, or does it mean that the two machines do not have XMS specification extended memory? If the latter is the case, what can I do with the 384Kb of extended memory that I have? I mainly use WordPerfect running through Power Menu.

Jeff Cole

One of the great weaknesses of the 80286 processor is that, although it is capable of directly addressing 16Mb of memory, it can only perform this feat while running in a mode that is incompatible with the wide range of DOS software available. Applications can switch into and out of the so-called 'protected' mode, which makes this memory accessible, but this is time consuming, and limits the memory to storing data — programs cannot run from extended memory, unless they are specifically designed to do so — like OS/2 applications, for example.

In its DOS-compatible mode, the '286 can only address 1Mb of memory in total, made up of 640Kb of RAM, and 384Kb of miscellaneous bits and pieces, such as BIOS ROMs. The extra 384Kb of RAM in your system lies above the 1Mb limit and is

inaccessible in the processor's real mode, (excepting the first 64Kb, which we'll get to in a minute).

However, the 384Kb space between the end of conventional memory and the start of extended memory is rarely fully populated, and with a bit of clever bank-switching, it is possible to fill these gaps with RAM, so that device drivers and other resident software can be loaded into memory, without impinging on the valuable conventional RAM below 640Kb.

The '386 chip has built-in bank-switching capability as part of its memory management unit, and with drivers such as Quarterdeck's QEMM, resident programs can be loaded into these vacant locations, by sacrificing some extended memory space, and switching the memory into these vacant spaces. Unfortunately, the '286 can't do this without outside help.

This help comes in the form of LIM EMS 4.0 — the expanded memory specification developed by Microsoft, Intel, Lotus, and a few others, to address the memory needs of 8088 and '286 processor-based machines. LIM 4.0 is a memory management scheme not unlike the one used internally in the '386 (albeit much less flexible), implemented in external hardware, either on a memory expansion board, or on the motherboard of more recent '286-based machines.

One of the things that LIM 4.0 can do is to put memory in the unused gaps between 640Kb and 1Mb, so that it falls within the address space of the processor in real mode, and can thus be used for device drivers and the like.

QRAM works with LIM 4.0 memory to achieve this end, by actually taking care of the loading of various bits of resident software into these high memory areas. However, it relies on the EMS hardware and software driver to already be in place to do this — it won't work on its own. Also, the version

161
SEPT 1991

of EMS must be 4.0 or later — previous versions did not have the flexibility to do the things required of it by QRAM.

I mentioned before that programs running on a '286 can actually address 64Kb of RAM over and above the 1Mb address space of the 8086 which it is nominally emulating. The reasons for this are related to the segmented nature of the Intel processors' address space, and I won't go into detail here. Suffice to say, some software can make use of this extra space, which is the first 64Kb of extended memory in any system, more readily than it can the rest of the extended memory space.

This special part of extended memory is known as the high memory area (HMA), and the Microsoft extended memory specification (XMS) defines a standard way of allocating this memory, as well as the rest of the extended memory in the system. Microsoft's himem.sys driver implements the XMS specification, as does Quarterdeck's qext.sys, for the '286.

Unfortunately, the HMA cannot be used as readily for loading resident software into as the aforementioned high memory blocks, but some programs can make use of the HMA to free up conventional memory for other things.

The best use that I know of for the HMA is with DOS 5.0, which can load part of its kernel (the central part of code which forms the operating system proper), into the HMA, leaving something in the order of 620Kb of the total 640Kb of conventional memory free for applications and data. This is achieved by putting the line:

DOS=HIGH

somewhere in your config.sys file. Note that this only works with DOS 5.0 — this is probably the biggest improvement of this new version

over previous ones. In your case, you really have two options — one is to go out and get an EMS board for each of your machines, and get the most out of QRAM, or get a couple of copies of DOS 5.0, and perhaps a disk caching program to use the rest of your extended memory. Sure, you won't be able to load TSRs into high memory as you could with the extra expanded memory, but it is probably the best use of the memory already in your system.

What's hiding in the shadows

Firstly, let me complement you on the interesting and varied format of your magazine. Your publication is one of the very few not filled with page after page of articles on how to use Windows.

I am currently using a 286-16 AT that has an AMI BIOS in it. In the CMOS setup program it gives you the option of enabling the shadow RAM for the system and video BIOSes. I tried both options and found that the system performance improved with both shadow RAMs enabled. I have not been able to find any documentation on how shadow RAM works. Can you enlighten me, and tell me when and how to use it?

Secondly, I have used several systems that give you the option of zero or one wait states. What criteria is used to determine which option to use?

B.J. McFadyen

We're glad to hear you like the format of the magazine, and are always eager to hear feedback on how we can best serve the PC community. While we believe that Windows 3.0 is an important product (why else would we give it the 1991 Software Product of the Year award?), we realise that it is not the solution to everybody's computing needs. There is a large

number of machines out there that is not capable of running Windows at a usable speed, and a large number of applications which gain no advantage from running under Windows. We're not ignoring the Windows market, but we are also committed to meeting the needs of the PC community at large, the majority of which is not using Windows as much as Microsoft would like it to be.

Now to answer your questions. Shadow RAM is a common work-around to solve a speed limitation with the type of memory chips used to store the system BIOS, and its various extensions, such as the video BIOS found on EGA and higher-order video cards. This software is stored in non-volatile read only memory (ROM) chips, which means that it is always present as soon as the computer is turned on.

Unfortunately, while processors and RAM chips have steadily increased in speed, ROM chips haven't, and when the processor needs to access one of these chips, it has to sit back and twiddle its thumbs for awhile to allow the memory chip to get the data ready. As processors get faster and faster, this waiting time (called wait states), becomes more and more significant. The BIOS routines are the lowest level of programming between the user and the CPU, and are called very frequently to perform virtually any operation which interacts with the computer's peripherals.

To avoid all this waiting around, many motherboards have the capability of implementing shadow RAM, which involves using much faster RAM to store the BIOS data instead of ROM. Since RAM loses its contents when power is removed, it needs to be copied from the ROM chips during the boot-up sequence. So the ROM chips are still there, but are only accessed once, very early in the start-up procedure.

After they have been copied to the shadow RAM, the shadow RAM is switched in to overlay the ROMs, so that all further accesses to the BIOS are made to the copy held in RAM. This is entirely transparent to the software running on the computer. The only difference is that BIOS routines execute much faster than they would from ROM.

This is the cause of the speed improvement which you have observed. I am not aware of any software which is incompatible with BIOS shadowing, since it is essentially transparent to the rest of the system. Also, since many motherboard configurations allocate memory for BIOS shadowing, whether you choose to use it or not, I can see no reason at all to disable it at any time. If you do, you're only unnecessarily shackling the machine, and not getting any benefit in terms of additional RAM.

When reference is made to wait states in system set-up programs or switch settings, they are referring to the number of wait states that the RAM imposes on the processor when it tries to read from or write to it. This is the same situation as already discussed in relation to ROM chips, but the number of wait states for RAM chips is much lower than for ROMs, since these chips are faster.

However, just as with ROM chips, RAM wait states impose a speed penalty on the system, and it is therefore desirable to minimise them as much as possible, and there are several ways to do this. The obvious one is to use chips which are fast enough to run with no wait states, but these are expensive, and the extra speed is often not worth the extra cash.

A compromise solution along the same lines is to use caching, where a small amount of very fast RAM (usually somewhere between 32Kb and 128Kb), is used to store the most frequently accessed parts of the main memory, which can

then be much slower. Another solution is to interleave the banks of RAM, so that accesses to successive memory locations are made to alternate banks of memory. Each bank thus has time to 're-charge' between accesses.

However, back to your question — the best way to determine whether to use a 0 or 1 wait state, lacking any other information, is to try the fast setting (no wait states), and see if it works. If it doesn't fall over (usually indicated by a parity error), after a fair amount of use (say, a day or two), then it would be fairly safe to say that the machine can operate at this setting indefinitely. However, if it crashes with such an error, then you would be better off switching back to one wait state and leaving it there.

Allowing easy selection of the number of wait states in this way let motherboard manufacturers make a single board for two models — differing only in speed. Slower RAM chips could be used with one wait state, while purchasers who were willing to spend the extra money on faster memory could benefit with increased speed.

Another laptop scanner solution

I am writing in connection to the query in the July issue of "Tech Tips", concerning the use of a hand-held scanner in conjunction with a Toshiba T1600 laptop computer. You mentioned two possible solutions to this problem, and I would like to add a third, which has been advertised in some American computer magazines.

This is a device which connects the scanner to the parallel port of the laptop, thus obviating the need for a slot in the computer to suit the scanner interface card. The device apparently works with most of the commonly used hand-held scanners, and is made by Computer Aided Technology Inc., 10132 Monroe Drive,

Dallas, TX, 75229. Their phone number is (214) 350 0888 in the US.

I hope that this information is of use to your correspondent. If his scanner is supported, at about US\$100 (from memory), it would be a considerably cheaper option than the Toshiba Docking Adapter or the Wonunder single card adapter, although plainly less flexible than either, being limited to just the scanner.

Jerome Stern

Thank you for supplying that useful piece of information — of which I wasn't previously aware — it sounds like it could be the ideal solution to those who want to use a scanner with their laptop without sacrificing its portability, aside from the obvious cost advantages. I don't know of anybody importing the device into Australia, but if somebody reading this knows otherwise, please let me know.

Of course, this device just adds to the range of devices connectable to the parallel ports of laptops in the absence of an expansion bus. External floppy drives, hard disks, tape backup units, LAN adapters, and now scanners. And, of course, you can always connect up a printer if the port is not busy running one of these other devices.

Boot up problems

I am an avid reader of Your Computer, and although not presently a subscriber, I have been in the past. I work at an institution which does subscribe, however, and I always make sure I am the first to borrow the monthly copy from the college library when it arrives.

The problem I write about concerns an Auva '386SX system, which I purchased earlier this year from Rod Irving Electronics. The computer has an 80Mb hard disk and 2Mb of RAM. The system won't boot from the hard drive,

161
SEPT 1991

only from the A: drive, where I have the three DOS 4.0 system files as well as autoexec.bat and config.sys. Once booted, it runs all programs quite all right from the hard disk, including Windows 3.0 and many others. When I try to boot from the hard drive, it goes through the BIOS routine without problem, and performs the memory check, but then beeps and flashes the message, 'No ROM basic'.

When I took delivery of the system, it booted from the hard disk, but wouldn't run Lotus 1-2-3. I reset some of the XCMOS NEAT chipset parameters, like the number of wait states and clock speed, thinking it might be a system speed problem. That's when I started having problems booting from the hard disk. I have since set all the parameters back to their original values, but I am still having no luck with booting from the hard disk, or getting Lotus to run. With the latter, it just jams up.

Is it possible that the problem may be within the ROM? If so, why does all the other software work? Could the Lotus problem be something to do with the fact that the system has extended memory, not expanded? I would be most appreciative if you could enlighten me.

I've contacted the dealer, and they say I must ship it back to their workshop in Melbourne. It's not that simple from here in PNG, as I waited two months for the system to be shipped in the first place. Also, I am still able to use the system, even with these problems, and I feel it should be fixable from my end somehow.

John L. Olney

I doubt that the problem with booting from the hard disk is related to anything you did in the XCMOS setup program. The fact that the hard disk is accessible in the normal way after boot-up seems to indicate that the drive parameters in the CMOS RAM are correct. I have seen some BIOSes which prevent booting from the A: drive, so maybe there are some out there which can lock out the hard disk as well. However, I think this is unlikely to be the problem.

The cause of the problem, I suspect, is that the boot sector of the hard disk has been damaged in some way. This is the first place that the BIOS looks for executable code after the A: drive, and it normally expects to find a boot record there, which is created when the disk is formatted. If it can't find this, then the next step is to look for a Basic interpreter in ROM. Since non-IBM machines don't have this installed, the BIOS has run out of options, and terminates with an error message.

If the boot sector is intact, and one of the system files is missing, then the error message is usually something in the form of 'Non-system disk or disk error'. It is difficult to guess what may have caused the boot sector to be damaged — most software should never touch it. One possibility is that a virus has infected the disk and destroyed the boot sector.

DOS doesn't have any in-built way to repair boot sectors that I'm aware of, short of backing up the entire drive, and re-formatting it, (with the /s parameter, of course), then restoring the data. The Norton Disk Doctor program in the Norton Utilities suite can repair damaged boot records, amongst other things, and other utility programs may also have a similar product. Failing that, I think the simplest solution is to back up anything of value on your drive, and re-format it. It would probably be a good idea to run fdisk first, just in case the partition table has been corrupted as well.

The Lotus problem, I think, is caused by something completely separate. You haven't specified which version of Lotus you are running, but, from the details in your letter, it sounds like version 3.0. This version uses extended memory, but it performs its own memory management in this area, unlike something like Windows, which uses an external XMS driver, usually himem.sys.

If Lotus 3.0 is run while himem.sys is loaded, the two appear to conflict, usually hanging the machine. Since you run Windows 3.0, I'd say that this is the cause of the problem. To get around it, you either need a separate config.sys file for Lotus (without the himem.sys driver) and another for Windows (and everything else), or a menu system to allow you to select different config.sys options at boot-up, such as boot.sys, which I discussed in Tech Tips in the October, 1990, issue. This is shareware, and is available from many bulletin boards, or we can send you and anyone else who has a need for it, a copy on disk for \$5 (no cash, please), to cover post and packing.

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on.

The best Tech Tip published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is Tech Tips, Your Computer, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

October 1991

Noisy disk drive

I have an IBM AT-compatible which will soon be clocking up its second year of use. The system was supplied with a Tandon TM-262 20Mb hard disk (a 3.5-inch drive mounted in a 5.25-inch bracket), which has recently developed an annoying problem. The problem is not with the data-storing capabilities of the disk, but with the mechanism itself. From the moment it was switched on, it sounded more like a high-revving petrol engine than a computer.

When it first started making this noise, it came and went at random, but finally it developed to the stage of being constantly there. Although this did not affect the operation of the computer, the operator never had any peace.

I tried everything short of dismantling the drive, and nothing seemed to work, until I found the solution by accident. Running out of room on my desktop due to the introduction of various new peripherals, I decided to turn the system on its side and put it under the desk, and bingo — no noise. The 'humming' had stopped completely, and I can now use my system without developing a migraine. Why did turning the drive on its side stop the humming, and will the silence last?

Zane Dawson

It sounds like you have (or had) a problem with bearing noise. This is not just limited to disk drives — any machine with rapidly rotating parts (virtually anything with an electric motor in it) can exhibit noisy bearings if they are poorly lubricated, or badly worn.

By re-orienting the drive, the weight of the spindle assembly of the drive is now applied to a different part of the bearing, which is apparently smoother than the part of the bearing which was supporting the weight beforehand.

It's hard to say whether the peace is going to last; I guess it is a case of 'enjoy it while it lasts'. I wouldn't recommend lubricating

the bearings at all, lest some of the lubricant find its way into the inner workings of the drive, and onto the oxide surfaces of the disks.

If the noise returns, you could always go and buy a heap of new disk-hungry applications. You'll then have to buy a new (bigger) hard disk to fit them all on, and the problem will magically go away!

Lost clusters under WordStar

I have an XT clone with a 20Mb hard disk, running DOS 3.3 and WordStar 5.5. Almost every time I run chkdsk/f there are two or three lost clusters to be recovered. On numerous occasions I have converted these to files, identified them and then checked the originals, only to find that they are intact. Does this scenario sound normal to you?

Ken Macnaughton

Is it normal? No, but certainly not unusual. Lost clusters are a frequent occurrence on DOS machines, but by rights they should not occur. Clusters become 'lost' when they are marked as being used in the file allocation table (FAT), but there is no entry in any directory that points to the cluster(s).

Chkdsk finds lost clusters by attempting to reconcile the entries in all the directories on the disk with the clusters that are marked as being used. Any clusters which are not claimed by a directory entry are assumed to be unused, and are tagged as lost. Chkdsk then gives you the option of converting these lost clusters back into files.

Each file contains one chain of lost clusters — a collection of clusters whose FAT entries each point to the next one in the chain, in the same way as when clusters are joined together to make a file.

Lost clusters usually arise when the program is interrupted during file operations. DOS programs almost always work with files through file handles, which are temporary aliases for the files which the application wishes to read from or write to. When a file open request is issued to DOS, it gives a file handle to the application, which the program then uses in all future references to that file. There is a limit to the number of files that can be opened in this way, set by the 'files=xx' line in your config.sys file.

If an application is terminated before the file is closed, it is possible that the FAT and directory entries may not match, causing lost clusters. This can happen if the machine crashes in the middle of an application, or the power is turned off without exiting the application properly.

It can also happen if you do something silly, like I did with the Syncomp 386 that I do most of my 'homework' on. Under DESQview, I was downloading a file in one window, while doing some file management with Norton Commander in another. In my efforts to organise the masses of files that one accumulates by relentlessly dialling-up bulletin boards and electronic news services, I accidentally moved the file I was downloading to another directory.

After the initial panic subsided, I noticed that the download was still proceeding normally, so I decided to let it continue. When it had finished, the file which I had accidentally moved was, not surprisingly, empty. However, running chkdsk reported that I had a couple of hundred lost clusters in two chains. I let chkdsk convert these to files and, behold, one was the file I had just downloaded, which was subsequently uncompressed and read with no errors. The other was a log file

OCT 1991

OCT 1991

created by Telix (my favourite comms package by a light-year), which was probably lost when I turned the machine off without closing the log first.

There is always a reason for lost clusters, and with careful observation, you should be able to figure out where they are coming from, and eliminate the problem entirely.

Mouse problems under Windows

Your answer to Robert Summers' disappearing mouse problem ('Tech Tips' July 1991), nicely explains why non-Windows applications (Lotus 1-2-3, and XTree Gold, for example) cannot use a mouse whilst running in a Window.

However, the loss of the mouse pointer in such programs when running in full-screen mode, or after quitting Windows altogether, was not covered.

Windows 3.0 has one of the best installation routines in the business, but even it is not perfect. For instance, it will only directly recognise mouse drivers which are loaded via a 'device=mouse.sys' style entry in the config.sys file. Those loaded via autoexec.bat are only recognised to the point that the installation check list will include the mouse amongst the list of installed hardware. Most of the time!

If a statement is found in config.sys, Windows will usually replace it with the line:

```
device=c:\windows\system\mouse.sys /y
```

I have found this mouse driver to give the best results under Windows, and to work just fine outside of Windows as well. What's more, on the cheaper Genius mouse (the GM6) which normally requires a button to be held down during boot-up to emulate a Microsoft mouse properly, this works very well without holding down the button.

If a mouse entry is not found in either boot file, but the hardware list during installation is amended to include a

mouse, Windows will invoke its own mouse.drv driver. However, this is only active while a Windows application is running. If you open up 1-2-3 or XTree Gold, there will be no mouse pointer.

Strangely, I have had more occasions where Windows has missed a mouse installed via autoexec.bat (leading to this problem) with the genuine Microsoft mouse (especially earlier versions) than with other brands. Microsoft mouse drivers which have not been updated by Word 5.0a seem particularly prone.

The solution is to invoke mouse.sys as shown above, except that you won't find it in your Windows System directory. Fear not. On Windows disk 2 there is a program called expand.exe. Copy this file to your Windows System directory; this can be used to unzip the file mouse.sys, which is on disk 4.

Place disk 4 in your floppy drive, and enter the command:

```
expand a:\mouse.sys c:\windows\system
```

The file will be unzipped and placed in the correct directory, but with the old extension. Type 'rename mouse.sys mouse.sys' to change it to a working form.

All that remains is to edit config.sys to include the new mouse driver in the config.sys file, using the 'device=' line above. You can use any ASCII text editor to do this, including Windows' own Notepad or Sysedit. Incidentally, with some mouse drivers, video errors can occur when switching between Windows applications and non-Windows applications. Typically, when returning to the non-Windows application, the screen will be full of blue and white stripes. The only safe way out is to exit that application blind, if you can. Not so easy with Ventura!

Microsoft Word 5.0a is a terror for this — great how different applications from the same company work together, isn't it? Mac users must be green with envy at the fun we PC users have.

Whilst on the subject of Microsoft products not co-operating, I have encountered many instances where a factory-supplied PIF shows full-screen operation, but the application always starts in a window, usually with warnings

of dire consequences and a total reluctance to actually run, but without letting you close it either. A sort of Catch-22.

This seems to be caused by the 'Video Ports: High Graphics' being selected in the Advanced section, (while editing in Enhanced mode). This feature also dramatically slows screen re-draws on those applications which will run. I turn that feature off with no discernible loss in screen quality, but a huge improvement in speed, and the application then opens in a window or full-screen, as directed by the PIF. Setting the High Graphics option does not seem to cause any ill effects, and is in fact essential for many programs.

Meanwhile, we are still trying to find out why the mouse pointer will randomly disappear when changing screens, only to reappear if we reverse the process and then go forward again. Also on the hit-list is the question of why the mouse pointer will sometimes wriggle off to a corner while the mouse itself sits perfectly still. Is there an invisible cat pointer chasing it perhaps? Does anybody know?

*Laurence Glynn
Sage Advice*

Thanks for those handy tips. Fortunately for us, most mice now emulate the Microsoft mouse at the hardware level, allowing us to use the genuine MS driver with virtually any mouse. Unfortunately, some older rodents, such as early Mouse Systems optical models (including the Vision mouse), were not hardware-compatible with the MS mouse, and if you have one of these, there's no way that the Microsoft driver will recognise it.

The two problems which you describe in relation to video (random garbage in non-Windows applications, and the Monitor Ports problems), are somewhat related. The purpose of the Monitor Ports option in a program's PIF is to inform Windows that it should keep a close eye on what that application is doing to the video display — the High Graphics box applies to high-resolution graphics modes.

Tech Tip of the Month — Displaying text files

```
@echo off
if '%1'==' ' goto HELP
if '%1'=='?' goto HELP

echo 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 > test . aux

copy test . aux /B + %2.com
if not errorlevel 0 goto ERROR
del test . aux
:HELP
echo This batch file converts text files to COM files so that they can be read easily
echo To run, use the following format: TXT2COM textfile comfile
echo where textfile is the DOS text file to be made executable, and
echo comfile is the name of the newly created COM file. The extension .com
echo is automatically appended.
goto end
:ERROR
echo An error occurred!
```

I WANTED A FILE that would allow me to display text files without any hassles. So I developed the accompanying batch file which allows text files to display themselves. It eliminates the repetitive use of the 'type' command.

The batch file can be created using the DOS line editor, Edlin. The mumbo jumbo of ASCII characters in the fourth line is actually a little machine-language

program which allows the text file to be executed. To enter these characters, simply hold down the Alt key while punching in the decimal numbers on the number pad, releasing the Alt key between characters.

The ASCII codes for those 'funny' characters in the fourth line are —

190-23-1-172-138
208-246-208-52
229-116-7-180-2
205-33-235-241
144-205-32.

Note: the character '32' is actually a space — don't leave it out! The batch file merges these ASCII characters with the text file, allowing it to be executed. Note that the original file is not altered or deleted.

Justin Gilbert

Why use a backup program?

Could you please help me with a few queries that I have?

First of all, could you please give a statement of the particular advantages of using a backup program (I am aware of many commercial offerings around these days), over simply copying important files from a hard disk to floppies using the DOS copy command, or a utility program such as PC Shell from PC Tools?

I use the latter method rather than a specific backup program, and find this is quite OK for my needs. Am I missing something?

I have a 42Mb HDD and find that I am getting a bit cramped for space.

What are the pros and cons of getting a second such HDD and using, say, DOS 5.0 to assign the two drives as C: and D:, over getting a single 80- or 100Mb drive?

I have been told that most PC viruses attack the C: drive; therefore, some say it is good to partition one's hard drive so that the logical C: drive is, say, a mere 2Mb, containing important system files, with the rest of the hard drive as drive D:, containing the ordinary program files.

Presumably, if one is infected by a virus, only drive C: needs to be purged. I would like your comments on this; I would have thought it would be reasonably easy to defeat such a scheme.

My compliments on an informative, interesting magazine.

John Lee

If non-Windows screens appear garbled when you switch back to them, it is a sure indication that the application is doing something to the display which Windows isn't aware of. The solution is to turn on the appropriate Monitor Ports switch in the PIF editor, and the problem should go away.

The penalty, as you have observed, is speed — a necessary trade-off for the ability to run non-Windows applications under Windows.

As to why this is causing some of your applications to crash, I can't really offer any help based on the information you've given us. If you can send in the names and version numbers of the programs in question, we'll attempt to get to the bottom of the problem.

Thanks for the bouquet, John. Now, let's hope I can keep up the good work and come up with an informative and interesting response to your questions!

The answer to your first question is simple — convenience. You can back up an entire hard disk by using the DOS copy command if you like, but it can be a headache keeping track of what files you have already backed up, and what ones are left. Remembering which directory each file came from is also not a trivial task when it comes to restoring them, nor is remembering on which disk a particular file is kept, if you need to restore that particular file only.

A backup program takes care of all of this for you, so all you have to do is to feed in the disks one at a time. The files are usually stored in some proprietary format (together with the relevant directory information), so that you need to use the associated restore program if you want to restore any files. Most intelligent backup/restore utilities put a complete index of the backup on the first disk, so that in order to restore a single file, you need only put the first disk into the drive, so that the restore program can locate the file, and then the disk that the file is actually on.

Another advantage of backup programs is that they automatically split files that are too big for one disk, and utilise all the available space on each floppy (except the last one, of course), to minimise

the number of disks used. They often compress and uncompress the data on the fly as well.

The decision as to whether to add a second hard disk to your machine, or upgrade to a larger one, is, for most people, financial — the former is usually cheaper, while the latter leaves you with an unused hard drive, that you may or may not have use for. If you do have another machine that you can use your old hard disk in, then replacing the hard disk may be the way to go, otherwise, I'd just add another drive of the same or greater capacity as the present drive, depending on your space requirements, (and budget).

If you decide to add a second drive, make sure that it is the same type of drive as you currently have — that is MFM, RLL, IDE or what have you. If you don't know what type of drive you have, ask your PC dealer, or a computer guru. Most drive controllers come with a daisy-chain drive control cable, but if not, you'll have to get one of those as well. Unless the drive is an IDE type, you'll also need to get an additional data cable as well.

If you decide to replace your existing drive with a new one, then you have the option of either sticking with the same type of drive as you already have, or changing, perhaps, to an IDE drive. IDE drives have become very popular of late, and are now standard in virtually all new machines on the market. If your current drive is not an IDE,

and you want to upgrade to one, then you will need to buy a 'paddle card' as well, which usually incorporates the floppy controller, and perhaps some I/O ports, too.

With any new drive you get, be it an additional or replacement unit, make sure that its drive parameters are supported by your system's BIOS, or that it has a customisable setting that will allow you to enter the parameters that you need. Some IDE drives can work with incorrect parameters, using a technique called sector translation, but it usually incurs a speed penalty of about 20 or 30 per cent. If in doubt, check with the drive supplier, and make sure that he is prepared to give you a money-back guarantee if it doesn't work.

The version of DOS that you use is unimportant as far as the number of drives is concerned. However, if you want to address each drive as a single partition, rather than chopping it up into 32Mb or smaller chunks, then use DOS 4.01 or later. It is worthwhile going for DOS 5 for the better use that it makes of extended memory over previous versions.

While it is true some viruses only attack files on the C: drive, most are more intelligent than that, and I wouldn't bother with making the C: drive really small, as the number of viruses that this would protect against is minimal. If you are worried about viruses, then get hold of one of the many good virus scanning programs available.

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on.

The best *Tech Tip* published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is *Tech Tips*, *Your Computer*, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

November 1991

Don't forget the '286

I feel Your Computer did Jeff Cole ('Tech Tips', September 1991), a disservice in your advice regarding his, and untold others', extended memory and QRAM problem. As 'home computer' users, we work within tighter budgets than corporations. After leaving the kids as down-payment for our 'modest' XT or AT, with a 'modest' 40Mb hard drive, a 'modest' VGA, a 'modest' 9-pin printer, and it would not be complete without a 'huge' 1Mb of RAM on board, sir', we probably have to mortgage the pet budgie to pay the Bankcard bill when it comes in. The computer industry's joke of the decade.

Ninety per cent of all magazines' editorials and reviews are too centred around '386 and '486 machines, with their 4Mb or 8Mb on board, the latest \$4000 software, laser printers, and so on. And if someone writes in with a problem, the best advice that can be given is to go and buy some more hardware and software. I read this as 'go and buy some more problems'. Anyway, I've had my whinge, now on to help Jeff.

I assume his Samsung is fitted with the NEAT chips, which allow Samsung owners to turn the BIOS shadow RAM on, which speeds things up a little. Reboot with the big red switch, and you will see some numbers rattle around (the internal RAM tests, or POST) — something like:

640K Base Memory, 00384Kb Extended
Strike the F2 key to run the SETUP utility

Before this test finishes, press the F2 key a few times to enter the machine's configuration program. Look down the column for the entry saying 'Extended memory: 384Kb', and change that entry to 0Kb. That's right, zero.

Press the 'Page Up' key, and a second setup screen will show the shadow BIOS setup. If all of these are 'Disabled', then you will need to 'Enable' some of them, or on boot-up it will report a 'Configuration Error'. (This also speeds up the machine's operation a bit — Ed).

I enabled System and Video BIOS

ROM shadowing, leaving everything else alone. Press the escape key to reboot, and the new POST check will show:

Extended Memory: 384KB

If you have as the first line in your config.sys file:

640KB Base Memory, 00000KB Extended

then you should be greeted with a message that says —

**High RAM created: 128K (B000-B7FF,
C800-DFF)**

I assume that you will write down the original settings before you change them, in case you stuff things up. (Who, me? Noooo!) Any rubbish in a config.sys file can lock up the machine, and QRAM's 'optimise' is unforgiving of a misplaced end-of-file marker. So, fellow experimenters, have a boot disk (or three), handy!

Peter R Cullen

Well, that just goes to show, you can't please everybody all of the time. I can't argue that we don't tend to concentrate on '386 and '486 machines, and at times the software which we review is rather expensive, (although very little comes anywhere near \$4000). The reason is simple — most of the new machines coming out are based on these chips. I have not seen a new XT released for quite some time, and the number of new ATs has dwindled to a trickle.

As a computer magazine, we try to keep our readers up to date with the latest advances in computer technology, particularly as they apply to the PC marketplace. With '386SX machines now the entry-level for new computers, it follows that these are the sort of machines which we look at when doing product reviews. You may notice that we have had very little in the way of reviews of '486DX machines, not because these machines are

hard to get hold of (we currently have a Protech '486 on loan which is doing sterling service as our office file server), but because we recognise that these machines are far more powerful than most of our readers will need in the near future.

However, we do recognise that there are a great number of XT and AT machines out there, which are giving excellent service, and have a good many years of work left in them. You may be surprised to learn that we still have a '286 here in our office, which we use as a front-end to our Harris typesetter. Once we dispense with the Harris and complete our cut-over to full desktop publishing, it will receive a hard disk transplant, and take over as file server from the Protech, (which Protech will hopefully let us hang onto until then).

Unfortunately, there are limitations to the '286 processor, and you can only push it so far before they become obvious. The biggest of these is the lack of compatibility between its protected mode (where it can address the full 16Mb of memory available to it), and real mode, where it is little more than a fast 8088, complete with its 1Mb memory limitation.

The NEAT (new enhanced AT), chipset breathed a little more life into the '286, by providing some of the memory management features which the processor itself lacked. Included in this are high RAM and EMS 4.0 support, which previously was only possible with special hardware, or a '386 processor.

If, as you say, Mr. Cole's computer has the NEAT chipset somewhere within its inner workings, then your suggestion will be of value. However, for the vast majority of NEAT-less '286s out there, there is little that QRAM, or any other software can do to alleviate the RAM cram. This is why we find ourselves recommending '386 machines as the entry-level more and more, as new software

NOV 1991

NOV 1991

gets ever more memory-hungry. The most significant side-effect of disabling the extended memory is that you won't be able to use DOS 5's new 'dos=high' feature. However, if you have a lot of TSRs loaded, and can put these into high RAM with the aid of QRAM, this will probably show a greater saving of RAM than DOS offers anyway.

Peter's point about having a boot disk handy cannot be overstated. Before attempting any changes to the config.sys or autoexec.bat file on any computer, make sure you have a reliable boot disk handy. If you don't, format one before you start, and copy your current autoexec.bat and config.sys files to it, before making any changes. And if you are going to play with the CMOS setup, make a note of those settings, too. It's a good idea to keep a record of your normal CMOS settings anyway, in case you suffer a battery failure.

Printer control codes

I have just finished Bruce Iliff's illuminating article about 'Using screen codes with your word processor' in the August issue of Your Computer, and couldn't wait to try it out on my IBM-compatible XT. Having consulted my printer manual (the computer manual and DOS manual were no help), to familiarise myself with the characters I could expect, I proceeded.

However, although I was able to reproduce all the characters from 32 onwards, I was unable to create the first 31 characters, despite using other methods as suggested by Bruce. Seeing as I had no trouble with the other characters, I find it hard to believe that either the software (WordStar 4), or the printer (Star NX1000), could be at fault, but I am at a loss to explain this course of events, and would like to hear from someone at YC or one of its many readers.

On another note entirely, I would like to ask if there is a legal way of altering exe files. I ask this due to the state of some of my software, that I received at the time of purchase of my computer several years ago. It was only recently that I became interested in several of them, and found to my dismay that, although the programs fired up, they did not work. In most cases there was little or no documentation available, or for that matter, any contactable person listed.

Hence my original question, to enable me to find out in the first instance how these programs work, and secondly, to make the appropriate changes.

Just a final note in closing, concerning the excellent work that everyone does at YC. I can only echo one reader's comments about the good advice and sound guidance that you provide.

Tony Beilhanz

Thanks for your positive comments about the magazine, and we're glad you find the magazine useful. Not all our advice meets with everybody's approval (see the item 'Don't forget the '286'), but we try to present as unbiased an opinion as possible in advising our readers.

The reason that the codes below 32 decimal act differently from the others is that these codes perform special control functions in the receiving device — in this case, your printer. For this reason, they are called 'control codes', and it is these codes that give the 'control' key on the keyboard its name. Pressing this key in conjunction with one of the ordinary letter keys (or a few others), causes a special control code to be generated instead of the normal ASCII code for the letter key being pressed.

Some of the more common control codes can be generated with a single key, labelled to indicate the purpose of the control function — for example, carriage return is Ctrl-M, backspace is Ctrl-H, and escape is Ctrl-[, The printer is designed to respond to

these codes is a special way — these three examples signal the printer to perform a carriage return, backspace one character, or signify the start of a printer control sequence, respectively. Therefore, whenever the printer receives one of these codes, it does not print a character, but performs some other special function, according to the value of the code sent.

These codes also have similar meaning for serial video displays as well, and are necessary when computers have serial terminals as their main I/O device, rather than the memory-mapped displays with which we are familiar in the PC range. With a memory mapped display, the CPU has complete control of the display, and can display any character at any desired position on the screen, and doesn't need to reserve some of the codes for control functions.

Because of this, IBM was able to (and did), put patterns for these characters in the PC's character set, and then any software could use them on screen as desired. Unfortunately, since the CPU has to control a printer via a single 8-bit port, the control functions of these codes had to be retained for printing. The codes from 128 through 255 pose no such problem, since they have no special control functions in the ASCII specification. In fact, they are not even ASCII codes, since the ASCII table only has 128 characters. Furthermore, WordStar uses these characters for similar tasks in its own file format, so any time that a decimal 13 value is encountered, WordStar knows that it has reached the end of the line, for example.

So unfortunately, there is not much that I can offer in the way of help here. Whenever anybody departs from an established standard, in this case, ASCII, problems of this kind are bound to arise when trying to interconnect the device with another which does

Tech Tip of the Month — Standard forms with Ventura

LATELY, MY WIFE and I have been producing a series of reports for clients involving up to 20 worksheet printouts from VP-Planner (a Lotus 123 clone), and up to 60 graphs. The graphs have proved troublesome. Although VP-Planner allows direct printing of graphs (rather than using Lotus Printgraph), the result was rather poor, with each graph occupying a whole page in landscape mode. Even worse was the time spent at the computer, with each graph taking up to five minutes to produce and print. What we wanted was three graphs per page in portrait mode. There are a number of alternative spreadsheets, but we could find nothing which allowed us to batch print 60 graphs from 20 different worksheets in a readable form.

We solved the problem using Ventura Publisher, producing graphs of startling quality and batching the entire set of graphs into one unattended printout. As each worksheet file was filled with data, we saved each graph as a pic file with the /GS command onto a floppy disk. We used a standard set of names for our graphs NUM1_P, NUM1_D, NUM1_I and so on, (these names had meaning to us).

We then set up a Ventura Chapter which contained three frames per page, each one linked to one of the standard named graphs on the floppy. The process of creating frames can be speeded up using the 'copy frame' facility within Ventura. Ventura reads Lotus pic files, converting them first to gem files so we had to allow room on the floppy for twice the number of files. The graphs read into Ventura quickly and without need for further scaling. Ventura handles graphics so well that the result far surpassed anything VP-Planner could produce. The client's details were incorporated into a

header and the whole lot printed out in one file.

Although setting up a Ventura Chapter containing some 20 pages and 60 frames took some time, we can now read any standard set of graphs from other clients by replacing the floppy in the drive. Ventura simply goes looking for a set of file names, finds them, and reads in the new graphs.

We did fall into one simple trap. We saved all our original pic files to the root directory of the floppy. It was quite comfortable with 60 files. By the time Ventura had converted all our pic files to gem files, we had 120 files in the root directory. In DOS versions up to 3.3, 360Kb disks are limited to 112 entries in the root directory. 1.2Mb disks may contain 224 entries. Ventura did not tell us that the root directory of our disk could not handle 120 files, it simply scrambled the directory. This was overcome by saving the graphs to separate directories on the floppy.

Terry Hinchcliffe

It's always interesting to see the solutions that people come up with to solve unique problems with off-the-shelf software. To many people, Ventura's chapter file format, with its absolute references to included files, is nothing but a hassle. Terry has managed to make this feature of Ventura work to his advantage, and developed a technique that would be useful to anybody who has the need to produce consistently-styled documents with different data.

Terry also earned a few brownie points by sending his tip (and the other one included in the main article body), on disk, which made my job easier!

However, we can't return disks, unless you include a self-addressed, stamped mailer.

conform to the standard, or differs from the standard is its own way.

As for wanting to alter executable files, I have to say that I think the legal aspects of this will be the least of your problems. Doing what you propose doing — taking an existing program, figuring out how it works, and then modifying it for your requirements, is far from being a trivial task. What you are talking about is known as reverse-engineering, and is the sort of thing that can keep a team of highly experienced programmers busy for months. And that's just for a single program.

Unless you have access to the source code of the program, it is

difficult to determine the structure of any program of significant size, let alone try to figure out why it doesn't do what you expect. I think the lack of documentation probably has a lot more to do with the trouble with the programs than any faults in the programs themselves.

I think you'd be better off trying to get hold of the documentation for the programs you have, or getting hold of new software (public domain, shareware, or commercial), which will probably have more features than the software that you are battling with now. If you start trying to disassemble those programs now, the PC will probably be but a dim memory by

the time you get them working.

Hard disk boot failures

Your September issue has a letter in 'Tech Tips' from John L. Olney, regarding a problem booting from the hard disk. I have experienced the 'No ROM Basic' error after running Fdisk on a hard disk and not setting the active partition. It is possible that the active partition was reset when John was modifying the NEAT chipset parameters, although I do not know whether the Fdisk parameters are stored in the CMOS or on disk. Could

NOV 1991

NOV 1991

you please forward this information to John — it could be worth checking.

Michael Moran

Thank you for your suggestion, although I don't think that it would help Mr. Olney in this case. Setting the parameters for the NEAT chipset is controlled by the built-in setup routine in the ROM chips themselves, and operates on a series of values stored in the battery-backed CMOS RAM on the motherboard.

Disk partition information is stored on the hard disk itself, in a special area at the start of the drive called the partition table. This information is accessed and modified in DOS using the Fdisk program, or one of a number of aftermarket utilities. The CMOS parameters are only significant as far as determining the physical characteristics of the drive itself. Once that has been established, the actual partition information is read from the partition table on the drive itself.

Normally, when running Fdisk, most users will select the default option for creating a primary DOS partition, to create as big a primary DOS partition as possible — up to 32Mb on DOS versions 3.3 and below, or the entire drive in later versions. If this option is selected, Fdisk will automatically mark the partition active. However, if you want to specify a different (smaller), size for the primary partition, then you must mark the primary partition active yourself, (still using Fdisk).

If you overlook this last step, the BIOS won't know which partition to boot from, even if there is only one on the drive. A drive can have up to four partitions, and any of them can be made active. Normally, this is only of interest for people with more than one operating system, each one with its own partition on the hard drive.

Changing the active partition allows the system to be booted up with your choice of operating system.

Your point is an important one to watch for, but I doubt that it was the cause of Mr. Olney's problem in this case. However, whenever you are re-partitioning a drive, make sure that you have the primary DOS partition marked as active, otherwise it will be inaccessible, even if you boot from a floppy.

Quick letterheads with Ventura

I often have a need to print on my own letterhead paper. Finding it annoying to have to unload my tractor feed paper and manually load a single sheet of letterhead, I solved the problem using Ventura Publisher.

My standard letterhead was designed using Ventura in the first place, placing a PC-Paintbrush image in a frame and adding text using the Box Text feature. Using this file, I printed the file to disk, using an Epson LQ width table. The resultant file was approximately 25Kb long and used the default extension .COO.

This file can be copied to the printer using the DOS copy command as follows:

```
COPY ARDHEAD.COO LPT1/B
```

The /B parameter ensures that the file is treated as a binary file. Since my wife is always forgetting DOS command syntax, I automated the process through a simple .bat file as follows:

```
COPY CON HEADER.BAT
ECHO OFF
ECHO Get the printer ready to print a
    letter head
ECHO Press CTRL-BREAK to cancel or
PAUSE
COPY ARDHEAD.COO LPT1/B
CTRL Z
```

We can now create a letterhead sheet of paper, 'on the fly', from either our

invoicing program, or our word processor, by shelling out to the system, (most modern programs have a facility for running a DOS shell without exiting). The resulting printout is just as Ventura itself would produce, but only takes a few seconds to create.

Terry Hinchcliffe

This is a tip that will be of interest to anyone using a computer to run a small business. Terry actually sent up two tips, but we felt the other one was especially ingenious, so we made it this month's Tech Tip of the Month.

New Boot.sys

Some time ago in 'Tech Tips', I made mention of a great little utility, called Boot.sys. This is a menu system that is invoked from the config.sys file, much like any other driver. What Boot.sys allows you to do is to determine which drivers you want loaded on boot-up — a feat not possible with conventional menu systems, since by the time they execute, all the drivers in config.sys have been loaded, and no others can be loaded (unless they have a com or exe form as well as the device driver form).

Boot.sys is now in version 1.42, and has been upgraded to allow it to take advantage of DOS 5 (although previous versions worked fine with this version of DOS). Boot.sys is designed for people who are forever altering their config.sys file to load different (often incompatible), drivers at boot-up. Boot.sys allows many different configurations to be selected at boot-up, without having multiple config.sys files, or having to edit entries in config.sys itself.

In fact, with Boot.sys, you can

New Boot.sys

NOV 1991

```

FILES=50
buffers=30
LASTDRIVE=Z
shell=c:\DOS\command.com c:\DOS /e:256 /p
device=c:\dev\BOOT.SYS /D1 /T8 /U^
device=top
device=top ^Boot.sys ^Menu
device=top ^Select one of the following options:
device=top
device=boot.1 ^D^E^S^Qview 386
DEVICE=C:\UTILS\DV.NEW\QEMM386.SYS NOSORT RAM=C800-DFFF
RAM=B000-B7FF
RAM=F500-F7FF
DEVICE=C:\UTILS\DV.NEW\LOADHI.SYS /b C:\DOS\ansi.sys
device=c:\UTILS\DV.NEW\LOADHI SYS /b C:\dev\mouse.sys
device=boot.2 ^Windows 3.0
device=c:\utils\dv.new\qemm386.sys NOSORT RAM=C800-DFFF RAM=B000-B7FF
RAM=F500-F7FF
DEVICE=C:\UTILS\DV.NEW\LOADHI.SYS /b C:\DOS\ANSI.SYS
device=c:\UTILS\DV.NEW\LOADHI.SYS /b C:\dev\mouse.sys
rem device=c:\UTILS\DV.NEW\LOADHI.SYS /b C:\windows\smartdrv.sys 1024 512
device=boot.3 ^D^O^S (normal, with ^T^S^Rs)
device=c:\utils\DV.NEW\qemm386.sys NOSORT RAM=C800-DFFF RAM=B000-B7FF
RAM=F500-F7FF
device=c:\utils\DV.new\loadhi.sys /b c:\DOS\ansi.sys
device=c:\utils\DV.new\loadhi.sys /b c:\dev\mouse.sys
device=boot.4 ^D^O^S - ^No ^T^S^Rs
device=c:\utils\DV.new\qemm386.sys NOSORT RAM=C800-DFFF RAM=B000-B7FF
RAM=F500-F7FF
device=c:\utils\DV.new\loadhi.sys /b c:\DOS\ansi.sys
device=c:\utils\DV.new\loadhi.sys /b c:\dev\mouse.sys
device=boot.5 ^Plain Old Dos - No XMM or EMM
device=c:\DOS\ansi.sys
rem device=c:\dev\mouse.sys
device=BOOT.END
dos=high

```

have up to 25 different menus, with up to nine options in each. This is much more than anybody in their right mind would want to go through every time the system is booted, but it's nice to know the capability is there, just in case.

So that the computer doesn't

sit there looking stupid if you don't press a key when prompted, you can set up a default option, to be selected after a pre-set time-out period has elapsed. Other features include the ability to re-define Ctrl-Alt-Del to perform a cold boot (instead of the more usual warm

boot), or to disable it entirely. There is also pause.sys — which allows you to pause the execution of the config.sys file for debugging purposes.

Boot.sys is distributed as shareware, and our copy came from Budgetware, (02) 519 4233.

December 1991

Windows files

I thought some readers may be able to benefit from a most perplexing experience I had recently. Basically, Windows went berserk! It refused to load any DOS applications, giving an 'Application Execution Error', proclaiming 'Cannot find file, check to ensure the path and filename are correct. It would get this error even if I used the File Manager to pick the file to run from a list displayed by Windows itself. I also noticed that several applications were using incorrect fonts — Turbo Pascal for Windows looked bizarre, and my background desktop bitmap was not being loaded.

Surprisingly enough (it was at the time, but not once the explanation became known), there didn't seem to be too many problems in Real Mode. I spent a couple of hours checking to see if I had deleted or corrupted any essential Windows files, paying particular attention to win.ini, which I had recently been cleaning up. I was just about to give up, when I noticed a brief error message flash past as the computer was rebooting — 'bad command in line 1 of config.sys'.

A quick check revealed that when editing config.sys, I had inadvertently put a square bracket in front of 'files=30' in line 1. With DOS ignoring this line, the default maximum number of files (8) is used. Problem solved, but to me, there are two points to be addressed.

Firstly, Windows should provide some meaningful error or diagnostic messages for this sort of problem. Previously Windows has refused to do things for me without giving any error messages, (this was in relation to problems with the 'load=' and 'run=' lines in win.ini).

The second point is that errors in autoexec.bat and config.sys should be highlighted more. Perhaps a 'Hit enter to continue' message. Most people start up a number of utilities in these files, quickly scrolling any error messages off the screen.

On another completely unrelated topic, I have a LightScan 400J scanner, which I usually leave connected to my computer. Unfortunately, DOS startup

activates the light in the scanner, and annoyingly leaves it on. I developed a small program to get around this problem, which I call from autoexec.bat. It can be entered using the Debug assemble option, or the hex values loaded using a small Basic loader or hex editor.

MOV	DX, 027a	(BA 7A 02)
MOV	AL, 00	(B0 00 00)
OUT	DX, AL	(EE)
MOV	AX, 4C00	(B8 00 4C)
INT	21	

If you want to turn the light back on (perhaps to send morse code to yourself), sending a 1 to port 27A turns the light back on.

David Johnson

Unless you are aware of the problem, it is all too easy to let error messages in config.sys go by unnoticed, especially if you have a rather busy config.sys and autoexec.bat file. One way around the problem is to put a 'pause' statement in the first line of autoexec.bat, so that you can read the messages that DOS puts on the screen while it is interpreting config.sys.

The trouble with this is that if some of the drivers that you load in config.sys display long-winded messages or, worse still, clear the screen at some point, then the message might already have disappeared before DOS gets to autoexec.bat. The latest version of the Boot.sys utility (see last month's 'Tech Tips'), included a handy little file called pause.sys. By putting the line 'device=pause.sys' in your config.sys file, you can make the display pause at any point so you can read the on-screen messages.

Boot.sys is available from many bulletin boards, or from Budgetware, (02) 519 4233. You still have to suspect a problem in config.sys in the first place, but once you know of the potential for

problems, a well-placed pause.sys will help you pinpoint it.

Corrupted boot disk

I recently bought a second-hand IBM clone computer with a 40Mb hard disk, with a few programs installed. My problem is that the 'system disk' is faulty, and the computer will not boot from the A: drive. I would like to make a few system disks to run as a front end to some minor programs, but I can't as a result of the inability to boot from the A: drive.

I don't have any other problems with the computer; it's only 10 months old and runs ideally for my requirements. I have tried various ways to copy DOS from the hard disk back onto a floppy, but with no success. I can't get at the hidden files. I'm running DOS 3.3.

There is plenty of documentation to copy DOS from floppy-to-floppy, and floppy-to-hard disk, but none seem to work from hard disk-to-floppy. Can you help me?

John Cole

Provided you have a working copy of DOS, either on a floppy or hard disk, there are two ways to create a system disk or boot disk. One is to format the disk from scratch, appending the '/s' parameter. That is, type 'format a: /s'. This will format the disk, and then copy the hidden files to the newly-formatted disk.

The other way, if you already have a formatted disk, is to put it in the A: drive, and type 'sys a:'. This won't format the disk, just copy the system files to it. This will fail if there are already files on the disk where the system files need to go, so it's best done with a blank disk. Also, with your version of DOS, it won't copy the command.com file to the disk, which is necessary in order to boot from it, so you'll have to do this manually.

If I were you, I'd create a few system disks as soon as possible, and keep them in a safe place. If your hard disk ever becomes corrupted, or the system files are accidentally deleted, you won't be able to start the machine at all.

Where's the cat?

Your reply to Lawrence Glynn did not answer all of his questions. I'm also having an intermittent problem with a mouse on one program, where the symptoms are exactly as described by Mr. Glynn. It is to do with the mouse pointer 'wriggling off to a corner', while the mouse is perfectly still. Can you offer any advice on this point?

By the way, I think you do a pretty good job with the magazine, spelling mistakes included. I do quite a bit of proof reading at work, and know how hard it is to try and maintain a perfect record.

Ray Villarroja

Without knowing more details about the software and hardware you are using, it is difficult to know where to start tracking the problem down. It seems as if the mouse is not quite compatible with the mouse driver you are using it with, or the driver is not compatible with the application.

The worst offenders are those applications which require you to use a mouse driver built-in to the application, rather than the one supplied with the mouse. Why programmers do stupid things like this I'll never know, but it does nothing to promote compatibility between their software and users' hardware. Driving the mouse is the job of the operating system, not the application, and the sooner application programmers realise this, the better. Admittedly, things are better than they were when the

mouse and the PC first met, but there are still far too many applications that insist on using their own mouse driver, rather than one built-in to the operating system, (as when you install a mouse driver in config.sys).

Applications with built-in mouse drivers generally work satisfactorily if the mouse you are using it with is on their list of 'compatible' devices, but if you're using a 'clone' mouse, problems often arise. The only general suggestion I can offer here is to try all the available options, in the hope of finding one that works properly with the application.

Thank you for your positive comments about the magazine — it takes a lot of effort to check all the material for accuracy, both technically and grammatically, and no matter how hard we try, little mistakes always slip through. It's nice to see that some people notice what we get right, rather than dwelling on what we get wrong, as is so often the case.

Lost clusters and bad sectors

In 'Tech Tips', October '91, both you and your correspondent Ken Macnaughton spoke glibly, it seemed to me, about converting lost clusters to files. When I run chkdsk/f on my C: drive, it blandly informs me that I have 40,960 bytes in bad sectors, but does not offer to do anything about it. I know it's only 0.2 per cent of my 20Mb, but our power supply is liable to fail at any time, sometimes several times on the same day, so my bad sectors are going to grow, and I'd like to keep them in check.

Incidentally, the D: drive has 24,576 bytes in bad sectors. As I never, as far as I know, use D:, how did it come to have open files when the power failed?

Celsus Clark

In the installment of 'Tech Tips' to which you refer, the discussion referred to lost clusters. Your problem — bad sectors — is completely different, although the effect of both is to rob you of hard disk space. However, unlike lost clusters, which are a temporary aberration, bad sectors are usually more permanent.

Lost clusters are the result of a logical inconsistency in the data on the disk, where disk clusters are not in use by any file, but are marked as being in use in the file allocation table (FAT). When looking for space to allocate to new files, DOS looks at the FAT to find the next free cluster. It will skip over any clusters marked as being already allocated, whether or not this marking is accurate. If the amount of lost clusters is allowed to increase unchecked, they will eventually occupy all the free space on the disk, and you won't be able to write any new files to the disk.

The DOS chkdsk program checks the consistency of the FAT with the disk's directory structure, and optionally frees up such lost clusters, (when the /f parameter is specified). If you answer 'yes' to the question about converting them to files, you can then peruse them to see if there was any valuable information in them.

Bad sectors, on the other hand, result from hardware problems in the drive itself. The two most common causes of bad sectors are imperfections in the magnetic oxide coating of the disk, or misalignment between the tracks on the disk and the heads. The former can be the result of either manufacturing imperfections, or may have been caused by a 'head crash', where the heads land on the surface of the disk where they shouldn't. Misalignment of the heads usually results from the continuous heating and cooling of the drive, and the resultant expansion and contraction of the drive components.

If the bad sectors are the result of an imperfection in the disk media, then there is nothing that you can do to recover them. If, however, it is an alignment problem, then performing a low-level format will recover them. If the number of bad sectors is slowly increasing, then it's most likely an alignment problem, and a low-level format should restore at least some of your bad sectors. If the number of bad sectors is constant, it is probably the result of a defect that has been there for some time, and there's nothing you can do about it. The bad sectors on your D: drive would most likely be the result of a surface imperfection, since you have never used it.

Since you have trouble with the reliability of your power supply, a few precautions will help ensure the maximum life from your hard disk. An autopark routine (available from many BBS' and user groups) periodically parks the heads of the drive over the designated 'landing zone', so that when the power fails, they will land safely, (unless you have saved recently). Also, if the power fails when the heads have not been parked, avoid bumping the machine until power is restored and you can park the heads properly.

Most voice-coil drives have auto-parking heads, which retract as soon as the power fails, so if you have one of these, there will be no problem with power failures. If you have a stepper-motor drive, then it's a problem you'll have to live with, at least for the life of the drive, and just do your best to extend its life. When it does eventually fail (all drives do, at some time or another), get an auto-parking voice-coil drive, and then the only thing you'll have to worry about are the easily-fixed lost clusters, next time the power fails.

The best solution to your problem, is to buy an uninterruptible power supply, or UPS. These monitor the power line voltage, and take over the power supply to the computer when the power is about to fail. Lumen Electronics — (03) 792 4203 — has a neat unit which is ideal for PCs. A full review appeared in April 1991.

Tech Tip of the Month Coffee-flavoured disks

A CLIENT came to me with a problem recently, which I thought I'd share with your readers — they'd a cup of coffee commit suicide all over a floppy disk that contained some important information, so they thought they'd try the resuscitation procedure previously mentioned in this column.

They carefully removed the disk from its cover, washed it in clean water with a few drops of a wetting agent, and replaced it in a new cover. It looked as good as new, but had no luck getting back the data that was needed. It even appeared as though the disk was completely unformatted — even a program that reads absolute disk sectors could not find a track or sector to access.

Out of interest, he reformatted the disk, and it came up with no errors. It really looked like the washing he gave it removed all the data as well as the coffee.

I guessed that his procedure was correct up until the final, and crucial, step. On a double-sided floppy drive, the two read/write heads are not mounted exactly opposite each other — instead they are offset by a small amount, so that each head presses the disk against a soft pressure pad.

In addition, the sector indicator (the tiny hole near the centre of the disk) has a specific relationship to the start position of each track. And finally, the formatting information, and the data are written to the disk while it is rotating in a specific direction.

Each of these factors mean that a formatted disk has a definite top and bottom side, and will only work when inserted in the drive the right way up. If a disk is always kept in its cover, then the question of

top and bottom never arises, as the drive will not attempt to access the disk unless it is inserted correctly.

However, if you have removed the disk from its cover and replaced it upside-down, then it will be absolutely unintelligible to the PC until it is re-formatted. This sounds like what happened to the client's disk.

Many disks have a reinforcement ring around the centre hole, making the top and bottom side easily identifiable. If the disk you are trying to rescue from the suicidal coffee cup does not have such a ring, then carefully mark it adjacent to the centre hole with a (water-proof!), marker pen, so you can be sure of inserting it into its new cover the right way up.

Jeff Richards

Oops, I had intended to mention that the disk needed to be put in the right way up, but that point obviously got lost somewhere between brain and disk. You are quite right about the importance of there being a top and bottom, although I'm not sure about the offset heads and pressure pads, at least in PC drives. The only drives that I recall having that 'feature' were the Twiggy drives used in the original Apple Lisa, and these needed special diskettes with two head openings.

As for hub reinforcing rings — they are popular on double-density disks, but not so on high-density ones because of the distortion this causes in the surface of the disk when the two are glued together. Still, this is handy to know, as it makes distinguishing double- and high-density disks easy, and with double density disks, you can tell which way is up. If you remove a high-density disk from its jacket before you mark its top side, you can work out which way is correct by trial-and-error. If it doesn't work the first way, flip it over (in its jacket), and try again.

Scanner compatibility

Recently my grandparents went overseas to England and bought me a Geniscan GS-C105 from a company named CYCA Electronics. It arrived here a week ago, and I've been having problems since then.

Firstly, the scan card that plugs into the motherboard has a 16-bit interface. When I first put the card in and turned the computer on, the computer didn't do anything. The fan was going,— but the computer didn't even start the system checks; it did not beep or anything. I investigated a bit, and found that it worked if I took a jumper off to make it work as an 8-bit card. It didn't matter if it was installed in a 16- or 8-bit slot, as long as the jumper was off. I tried it in another computer, and it started up with the jumper in the 16-bit position.

My computer is a 16MHz '286, with 2Mb of memory, two floppy drives, a hard disk, a 512Kb Oak SuperVGA card, a SoundBlaster and a Vidi PC video digitiser. I also tried the scanner card with the SoundBlaster and Vidi PC cards removed, but it still didn't work. The other computer was exactly the same, but without the Vidi PC, and a different VGA card. The motherboard was also different.

My second problem is with the software that comes with it, called 'Maestro'. I can only get it to work with 16 colours in 640 x 480 mode. Although it doesn't have the Oak driver, I tried all the drivers that came with it, but none worked with more than 16 colours. It supports Tseng Labs VGA cards, so I borrowed a 1Mb Tseng Labs card to try it out. This time at least I got something on the screen in 256 colours, but it appeared squashed up and shaky in the top section of the screen. Somebody told me that this was because the frequency of the card has to be changed, but there was no DIP switches or jumpers to change anything on the card. Can you tell me what I can do about this problem, because with my Oak VGA card I can use 256 colours in 640 x 480 and 800 x 600 modes with Windows and the Vidi PC? Must I change the card, and if so, what type of card do I need?

Could you please give me a contact number of someone that deals with Geniscan products, if you can't resolve my problem? Could you also give me a name of a good memory simulator, and where I can obtain it? I have a public domain one, but I think it has been corrupted, because it doesn't do anything.

Tesaro Sandu

It sounds like the scanner card is conflicting with the video card in your computer, although without knowing the specific details of the components in your system, it's difficult to be more specific. You could try swapping the video card for the one in the other machine, and see if that works.

It's not unusual for cards to conflict while in 16-bit mode, and yet work perfectly well when forced into 8-bit mode. I know of several cases where 16-bit VGA cards didn't work in a system with IDE hard drives, yet would work perfectly well if the VGA card was switched to 8-bit mode.

Running the card in 8-bit mode may not be that much of a handicap — a scanner is not the sort of device you would use all the time — and at worst, all this will mean is that you may have to pass the scanner over the material to be scanned a little slower than you would in 16-bit mode.

As for your Oak VGA card not working with the software in SuperVGA modes, this is not all that surprising considering the lack of standardisation of video cards in modes other than those present in the original IBM article.

A few years ago, several manufacturers of high-performance video cards got together and formed VESA, the Video Electronics Standards Association, to draft a standard software interface for video cards which go beyond IBM's original implementation.

Unfortunately, because all of these manufacturers were already

committed to their own hardware implementation, coming to any sort of agreement at the hardware level was clearly out of the question, and the result was a standard device driver interface. It was up to each card manufacturer to provide a VESA driver for their particular card. All a software writer would have to do then was to provide an interface to the VESA driver, and all would be well.

Of course, nothing ever goes as smoothly as this in the computing industry, and we are left with a horrible muddle as far as SuperVGA support is concerned. In fact, the term SuperVGA seems to be the only thing that all software and hardware manufacturers seem to agree on! Some software still comes with a host of different drivers for different video cards, some (very little), comes with a VESA driver, usually in addition to the card-specific drivers. Support amongst the hardware manufacturers is little better, with more cards being sold without VESA drivers than with them.

Until this situation is resolved, there are going to be many, many, compatibility problems. The software and hardware manufacturers each say that the responsibility of providing drivers belongs to the other. Windows and OS/2 solve this problem to some extent, since you only need to provide a single driver to interface the operating system to the card, and all software works through the operating system. Of course, this is only applicable to Windows or OS/2 applications.

Since you have already tried all the existing drivers (often some manufacturers clone somebody else's card, but put their name on it, so the drivers are compatible even if the names are different), you only really have two options. You could try to convince either the software publisher, or the makers of the Oak card, to provide a suitable driver (good luck!), or get a new

video card that will work with the software.

The most likely cause of the problem you describe with the shaky picture with the Tseng Labs VGA card is that the Tseng Labs driver in the software may be incompatible with the card you are using. You are fortunate in knowing a dealer who is willing to let you test the card with your software before shelling out your hard-earned cash the industry is full of people who would rather just grab your money and run.

There are two common Tseng Labs chip sets in VGA cards the ET3000 and the ET4000. Some of their modes appear to be compatible between the two chips, but others are definitely not. Make sure that the driver is designed to work with the particular chip that is on your card. If the software doesn't specify which chip it works with, then I'd say it is most likely — designed to work with the ET3000, the older of the two chips.

It's also possible that the monitor you are using is not compatible with the SuperVGA modes offered by the Tseng Labs card, although the fact that you can run Windows in 800 x 600 mode with your other card makes this possibility rather unlikely.

The local distributor for Geniscan is Pactronics, (02) 748 4700.

As helpful as the support staff there are, keep in mind that you and your grandparents are guilty of 'parallel importing' the scanner and they are under no obligation to help you. (Have you thought of contacting the English supplier?)

As far as EMS emulators are concerned, the best '286 emulator we've seen is a product called SoftBytes.

It only emulates EMS 3.2 — emulating EMS 4.0 fully on a '286 is not possible without special hardware. Our copy came from Software Express — (03) 663 6580.

Disk densities revisited

In a recent computer magazine (I am not certain which one), one of the apparently knowledgeable staff claimed that Teac 1.2Mb floppy drives were capable of writing without errors to 360Kb diskettes because the Teac heads were slightly wider than most 1.2Mb heads.

At the end of last year, I purchased a '386SX system with both 1.2Mb and 360Kb floppy disk drives, on the understanding that both were mandatory to successfully write to each kind of diskette. I understood that while a 1.2Mb FDD head may correctly write to a blank (newly formatted), 360Kb diskette, when one tries to do the same to a 360Kb diskette which was full and has merely had all the files deleted, the narrower 1.2Mb head will not completely overwrite the wider 360Kb tracks.

At the time my dealer cheerfully sold both types of FDD to me. Recently he just as cheerfully admitted that in their office, they use a 1.2Mb FDD to write to 360Kb diskettes! If this can in fact be done safely using the Teac FDD (and perhaps others), is there any conclusive test you could suggest that would, without doubt, prove that the particular 1.2Mb FDD one had, could safely write to 360Kb diskettes? I suspect not, because if it could be done, wouldn't manufacturers openly tout their FDDs with such a capability?

John Lee

This is one of those issues which crops up from time to time, ever since IBM released the AT. The users' and technical manuals for this machine warn users that 360Kb diskettes written in its (1.2Mb) drive may not be readable in a 360Kb drive. This was true of the drive used in the IBM machine and also most clones.

Then, along came the Teac drives. I don't recall Teac ever making any official statement about disks written in their drives being compatible with 360Kb drives, but in my experience with many machines equipped with Teac drives, as well as that of many associates, 360Kb diskettes written in 1.2Mb Teac drives have been readable on any 360Kb drive I've tried them on.

Admittedly, I was a little skeptical myself when I heard that this was possible, but I have been unable to fault them. I don't know whether Teac drives have wider heads in them, or use a higher writing current, which swamps the signal on the unwritten half of the track, but it definitely does work.

Can I prove this statement? The only answer I can honestly give to that is, 'no'. To do so would require measuring equipment capable of measuring the signal-to-noise ratio of disks written on these drives, and comparing it to the same disks written on 360Kb drives. I don't have access to this sort of gear myself and I haven't not seen anything written by anybody who has performed such tests.

I don't think your dealer was taking you for a ride by selling you both a 1.2Mb and a 360Kb drive. I think he was just playing it safe.

If he had claimed that you could use the one drive for both disk densities, and you had subsequently found out otherwise, it would have left him in a difficult position.

I myself used to recommend that people get a 360Kb drive if they wanted to be certain of compatibility, although now I usually recommend people get a 1.2Mb Teac drive, and a 1.44Mb drive as the second drive.

January 1992

Is it a virus?

I have an IBM XT computer with a 21Mb hard disk, with DR DOS installed. My problem is that a bad sector on my hard drive is growing by about 10Kb every 12 hours of use. I have scanned the hard drive for a virus, but there is none on the hard disk, and I have also tried a disk doctor, but all that it tells me is that I have a bad sector on my hard disk.

I also have another problem, with my printer. Every time I have to use it I have to type in:

MODE LPT1:80,6

or it will print garbage characters.

Steven Cantle

This sounds like a classic head alignment problem. It arises because of the open-loop nature of the head placement mechanism in stepper-motor hard drives — there is no feedback to ensure that the heads are positioned exactly over the tracks on the surface of the disk. Wear on the drive can cause the position of the heads to drift slightly — enough to cause trouble reading data from the tracks in places. It is not uncommon for such bad spots to grow in size, even though there is nothing physically wrong with the drive.

The solution is to re-format the drive, although thanks to some utilities available today, this need not be the arduous task that it once was. However, before proceeding, it is a good idea to backup the drive anyway. That way, in the unlikely event that something goes wrong, you can still restore your data to the drive.

Spinrite is one of the most popular on-the-fly re-formatters, which reads a track into memory, re-formats the track, and then writes the data back out. Generally, nothing will go wrong, unless a

power failure occurs between when the track is re-formatted, and the data is written back to it. So it's best to play it safe and do a backup first. If you already have copies of all the data on the drive elsewhere, you can dispense with the backup step. For information on Spinrite, contact Mace Reel Software, on (07) 252 3992.

The other way, if you don't want to buy Spinrite, is to backup the entire drive, and then perform a low-level format of the disk the old way. With most XT clones, the method is to enter debug, and then type 'G=C800:5'. This will run the low-level formatter built into the ROM on the controller card. Genuine XTs need the IBM Diagnostics Diskette, or a third-party program like Ontrack's Disk Manager. AT machines usually come with this utility on disk.

Following this, run fdisk to re-partition the drive, and then format the drive, before restoring the data. I have performed this operation on many machines (before Spinrite came on the market), and in all cases it cured the 'growing bad sectors' problem.

As for your printer problem, it would appear that the printer is not initialising properly. The mode command which you have typed puts the printer into 80 column mode, with 6 lines per inch vertical spacing. Just put this in your autoexec.bat file, and it will initialise the printer each time you boot up without you ever having to worry about it again.

If anything can go wrong ...

The proliferation of IDE hard drives has added a new twist to the setup routine of many PC BIOS' — the user-definable drive type.

As anybody with an IBM AT or

clone knows, there are a number of important parameters about the machine stored in a little thing called a CMOS RAM, buried somewhere within the bowels of the computer. Most people only ever worry about them when they need to change one of the parameters, like after adding a new hard disk, or more memory.

The other time the CMOS RAM makes its presence known is when it is struck down by amnesia. Suddenly, the computer has no idea what size disk drives it has, and therefore can't access them. Oh, it can see that the drives are there, all right, and therefore knows something is amiss, but there is no way that it can figure out the size of the drive and reset itself. After all, if it could, there would be no need for the CMOS data in the first place.

The CMOS RAM is actually part of the clock/calendar chip which every AT (and '386, and '486, and ...) has on its motherboard. The clock chip has a little battery connected to it to keep it running when the mains power to the machine is turned off, and as well as keeping the clock running more-or-less on time, the battery stores the critical values in the CMOS RAM, ready for the next time the computer is turned on.

All it takes is a brief short-circuit between the wrong bits on the motherboard, and the power to the CMOS chip will be interrupted long enough for it to forget everything it ever knew. At least, that's what I think I did the other day when poking around inside my current home machine.

I was installing an NEC SCSI host adapter, so that I could check out a few CD-ROMs. Off came the blanking plate for a spare slot, in went the SCSI card, followed by the retaining screw. Put the cover back on, tightened up all the screws, and plugged everything back in.

I knew something was wrong the instant the RAM self-test

started — the little counter which checked off the memory in 64Kb blocks as they were tested was ticking over at only a quarter of its normal pace.

After the memory test had taken an eternity to complete, the two floppy drives buzzed in sequence, as they were tested. Then a pitiful beep from the speaker, accompanied by an equally pitiful 'drive not ready error' message.

Oh dear, something's really killed it this time. First suspect was the SCSI card — after all, that was the last thing I had added to the machine.

Down I went, on hands and knees, under the desk, out came all the cables from the back of the machine, and out came the machine from under the desk. Tower cases are a wonderful idea, until you have to work on them. Then you start wondering whether a desktop case was all that obtrusive after all.

Out came the SCSI card, and back went the box (sans cover, this time), and in went just enough cables to make it work — power, monitor signal, monitor power, and keyboard. A flick of the big grey switch (I really prefer big red switches; they have a purposeful look about them), and the machine sprang to life.

Or so I had hoped. Truth is, it did exactly as it did the time before — slow RAM test, and then hard drive error.

That ruled out the SCSI card — not surprisingly, since this card has been in more machines than the stoned virus, and the only one it ever failed to work in was a Wyse '386, in which every other card I tried failed also.

By now it was clear that my main work machine at home was seriously ill, and what's more, Jake was expecting my article the next morning.

Once more I hit Ctrl-Alt-Del, but this time I pressed the delete

key after the RAM test, which forced the AMI BIOS to enter its setup routine.

What I found there was not a pretty sight — everything had been zeroed. The time was just after midnight, on January the first, 1980 — the beginning of time as far as any PC BIOS is concerned.

What was worse, the entry for the first hard disk (the only one I have in this machine), said 'not installed'!

I uttered a few appropriate expletives, and set about the task of resetting all the values in the CMOS table.

The only things that remained intact were the settings for the two floppy drives — quite probably because the BIOS figured them out for itself when it was turned on. The hard disk was a type 47, I remembered — the AMI BIOS' number for a user-definable drive.

When type number 47 rolled around, my worst fears were realised — all the parameters for the hard drive had been reset to zero also. Not surprising, in retrospect. After all, they're stored in the same CMOS chip as all the other parameters.

The trouble was, I didn't have a clue what they were. The hard drive was a Fujitsu 2612, a drive I have encountered in many machines, but of course, I didn't have the parameters handy.

The moral of this little story should be obvious — know thy computer.

Keep a copy of all the CMOS RAM settings for your machine, in a safe place, in case this should ever happen to you. Make it part of the 'recovery kit' that every computer user should have — another essential component of which is a DOS system disk, complete with all essential device drivers, copies of your autoexec.bat and config.sys files, and even your favourite editor, in case your hard disk should lose its system files one day.

T1000SE won't boot

I have a Toshiba T1000SE laptop, which has served me well for over a year now. The other day, in my fiddling, I set it to boot from the RAM disk, rather than the usual ROM disk, and that's where I came unstuck. Before I had a chance to put the system files on the RAM disk, the machine had re-booted, and gave the familiar 'non-system disk' error message.

Also, when the RAM disk is selected as the boot drive, it doesn't check the floppy drive beforehand, so there was no way that I could boot from a floppy disk to rectify the problem. I eventually disconnected all the batteries (which means opening the machine up, something I would rather have avoided had the machine still been in warranty), and then re-booted. Is there some way of avoiding ripping the machine apart in these circumstances?

E. Chung

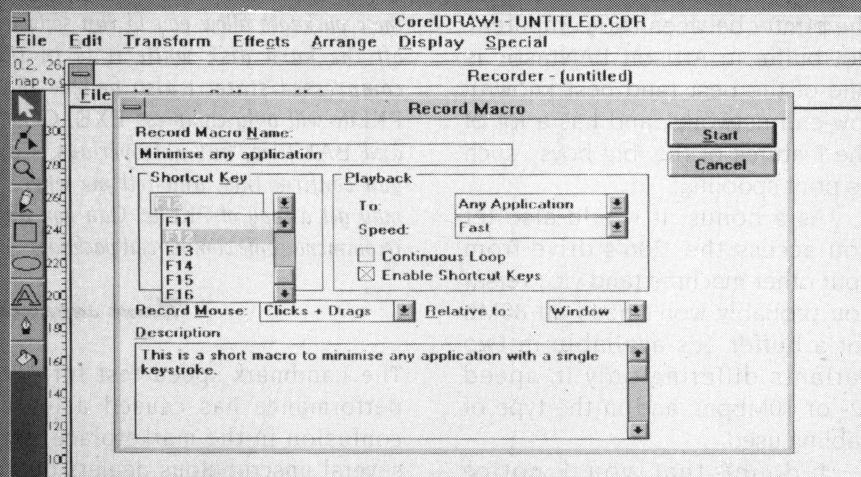
I myself have run into this problem a couple of times, and rang Toshiba's tech support people. It turns out there is a simple solution to the problem hold down the F1 key while turning the machine on. This overrides the settings determined with the setup program, and forces a default boot sequence, from the floppy drive, then the ROM disk.

QRAM and extended memory

I have just read your excellent response to Jeff Cole's enquiry in the September edition, which has cleared up a lot of confusion on the subject. I still have a query or two, though!

At the moment I also have a '286 (Imagineering Ultra '286 8/12MHz), with 1Mb of RAM and have been trying to access memory beyond the 640Kb

Tech Tip of the Month — Windows recorder



IF YOU haven't started to use Windows' Recorder, here's a helpful macro you might like to record to get started. It minimises a program to an icon with a single key. I have chosen F12, but any key or key combination is possible.

Once in Windows, launch one program from Program Manager, maximise it, and then switch windows back to Program Manager, making sure the program you left is still maximised. One way to do this is to hold down the Alt key, and at the same time, continually press the Tab key until Program Manager is highlighted. Once you have toggled to Program Manager, release the Alt key.

Double-click on Recorder, and you are ready to go. Click on Macro and then select Record. In the name window, enter a description of the macro, such as 'Minimise any application'.

Next, select the key that you have chosen to initiate the macro, and put it in the Shortcut Key window. Either enter the key from the keyboard, or click on your selection. Then select or deselect Ctrl, Shift or Alt, if desired.

The 'Record mouse' box should be left as 'Clicks + Drags', and leave 'Relative to' as 'Window'. However, 'Playback to', should be changed from 'Same application' to 'Any application'. Leave the speed as 'fast', 'continuous loop' unselected, and 'Enable Shortcut keys' as active. You are now ready to record the macro, so click on Start.

The Recorder will minimise, and you will be returned to the program you previously started you can now record the macro. Because not all

programs maximise to a full screen, you should not record mouse clicks, but rather minimise the program with the keyboard. To do this, press Alt-space, followed by 'n', for 'miNimise'. Your program will now be reduced to an icon, and Recorder will start flashing to indicate that it is recording.

Click on the recorder icon, select 'Save macro', and click on OK. Your macro has now been recorded.

Double click on Recorder, and save your macro in a file — main.rec is as good as any to start with. Each file can contain several macros, but you can only load one file into Recorder at one time.

For your new macro to be active on start-up, you must either have Recorder loaded as an icon, or maximised with the file loaded into the recorder. To load Recorder and your macro file automatically when Windows starts, edit win.ini, in the Windows directory. First make sure you have a backup copy in case of errors, then edit it in the same way as you would edit auto-exec.bat, or config.sys. The line beginning with 'load=' near the top of the file should be modified to contain Recorder and its path, followed by the file you want to load, thus:

```
load=c:\windows\recorder.exe main.rec
```

Once the changes have been made, you can restart Windows, and test your new macro.

PE Bailey

mark, but without success. I would like to access more memory, as I run Lotus 1-2-3 with the Allways, InWord, and Guardian add-ins, which take up lots of memory. I frequently see it bomb out with the memory error message when all the add-ins are attached.

Amongst my collection of bits and pieces, I have a 2Mb XT RAM board, plus LIM 4.0 software leftover from when I had an XT clone, and I also have the QRAM software.

The first thing I did with the '286 was to try QRAM, but got only the same message as Jeff Cole; the LIM 4.0 EMS driver was not installed then. If I install this driver before QRAM in the config.sys file, will this fire it up?

Will the XT RAM board work OK in the AT, or will I have to get an AT board. I looked inside the PC the other day, and note that the on-board RAM is not the same as on the old XT there are only two rows of chips with about five or six in each row. Can I get 2Mb on-board?

Lastly, what would you recommend as a good disk caching program?

David Peet

I can see no reason why an XT EMS RAM board would not work in an AT. Although it would operate slower than a specialised AT card, it would certainly be better than nothing.

One point to watch is that, although some boards come with EMS 4.0 software, full support of this standard requires special hardware, which some cards do not have on-board. These cut-down implementations may not allow QRAM to implement all of its features, but you should still get some advantage from QRAM. The best way to find out is to install the memory card, and then put the EMS driver and QRAM in your config.sys file (in that order), and see what happens. The Manifest utility supplied with QRAM will be of some help here.

Note that this will leave the extra 384Kb of memory on your

motherboard unused, but if you upgrade to DOS 5.0, you will be able to use 64Kb of this memory for DOS, leaving more memory below 640Kb for your applications. The rest of the memory could be used as a small RAM disk or cache.

I don't know whether you can get another megabyte of memory on your motherboard, not being familiar with that model of machine. However, unless the machine is based on the NEAT or SCAT chipsets, any additional memory thus installed would only be usable as extended memory, not EMS, as Lotus requires. However, you could use it for a RAM disk or disk cache, if you like.

I haven't tried a lot of disk caching programs myself, but often use the one which comes with the Norton Utilities, version 6.0. This can use extended or expanded memory, and a vast array of options such as delayed writes and Windows 3.0 compatibility, and I certainly have not had any trouble with it in the time I've been using it. For details, see John Hepworth's full review in the January 1992 issue.

Computer as a print spooler

I enjoy the magazine, and find trying to make my computer work good therapy. Can you advise me on how I can use my spare '286 machine (with an 80Mb hard disk), as a spooler, (or is that buffer)? I have to print some large files, and I am (typically), impatient to spend my time on the next 'main event'.

Clive Robertson

We're glad to hear you like the magazine. We hope to make it useful to most people most of the time. It sounds like you need some sort of network. Other similar

products tend to only let you share the printer between two computers, not buffer it. Artisoft LANtastic is one of the best (and best known), low-end networks, and has a lot of the features of the 'big boys', such as print spooling.

As a bonus, it would also let you access the '286's drive from your other machine (and vice versa) you probably won't need all 80Mb for a buffer. It's available in two variants differing only in speed (2- or 10Mbps), and in the type of cabling used.

I doubt that you'd notice much difference between the two versions (if any), for printing operations, and the limiting factor is usually the speed of the application doing the printing. Unless you plan to expand the network significantly in the future (say, to 10 machines or so), the 2Mbps should do fine for both file and printer sharing.

The Australian distributor for LANtastic is Digital Solutions, (07) 883 1851.

CPU's and speed

Can you please explain the difference between the Landmark speed on a computer, and the CPU speed. In other words, what is the use of having a 33MHz '386 when you may get a 25MHz turbo running at 33MHz, (or more, in one case I saw recently)?

Can you also explain whether or not the '386 will run software in the future. Having bought an XT less than two years ago and having it declared redundant, I do not want to spend thousands of dollars on buying a '386 and finding two years down the track that it will only run 'ancient' software.

The second part of my question is in reference to hard disk storage. I realise that there are some software packages around that allow you to increase the

storage of a hard drive. Can you tell me if these packages allow you to run software off the hard disk while it is in the compressed state. I also believe that PKLite will only compress EXE, COM, and BAT files, but not overlays. Can you confirm this, and tell me where I may get a copy of PKLite. Can you also recommend some commercial packages

Name supplied

The Landmark speed test for CPU performance has caused a lot of confusion in the marketplace, and several unscrupulous dealers quote this figure as if it were the true clock speed, which it isn't. The Landmark test was devised to compare the performance of the test machine to a known benchmark the IBM AT. Thus, a Landmark speed of 8MHz means that the machine is the same speed as an 8MHz IBM AT.

The reason that Landmark figures are often higher than the real CPU speed is due to wait states, or the lack thereof. The original AT had rather slow RAM chips, and to allow these to work with the (then) fast CPU, a wait state was inserted in each memory access, causing the machine to wait for an extra clock cycle before continuing.

As faster RAM chips became available, this wait state could be dispensed with, and the machine would run faster for a given clock speed. So, the Landmark speed for a 10MHz machine running without any wait states is about 13MHz. Sales people like to quote Landmark figures because they're higher than real clock speed.

So the point to watch is to make sure that when comparing machines, you compare Landmark figures to other Landmark figures, not the actual clock speed. This is the best way of quickly comparing CPU performance.

The '386 family of chips is going to be with us for quite a while yet, and new software is going to appear for this family for quite some time. The 8088 used in XT machines is quite old technology, and the only reason that people kept selling them was because

February 1992

most of the available software would run on them.

The '286 was not useful for much other than as a fast 8088, so it didn't stimulate much in the way of new software. It wasn't until the '386 arrived, with software that could take advantage of its powerful new features, that software manufacturers started producing software that would only run on the '386. This has happened quite recently, but was foreseen by industry observers for quite some time.

The next big step is likely to be towards RISC processors, but these are unlikely to have a major impact on the mainstream PC market until late in the decade. The '486 (and the future '586), are basically much-enhanced versions of the '386. The software that they run is the same as the '386 chip, only faster. So the '386 is going to be able to run all PC software for the foreseeable future.

One of the best on-the-fly data compression utilities we've seen lately is a program called SuperStor. It compresses part or all of your hard drive, and just uncompresses files as they are needed completely transparent to the operator. The slowdown in speed of the disk drive is barely noticeable in most cases, and it can achieve compression rates of about 2:1 with an average mix of data.

PKLite is a different story. Rather than being a device driver which intercepts disk accesses, it basically turns each executable file into a self-extracting archive. When it is run, PKLite loads the real program into memory, uncompressing as it goes, and then passes control to the program when it is uncompressed.

Since overlay files and data files are never actually 'run' as far as DOS is concerned (even though they might contain sub-routines which are subsequently called by the main program), there is no way for PKLite to work with these files. So if you want to compress everything on your drive, SuperStor is probably the way to go. SuperStor was reviewed in our October 1991 issue, and is available from Data Image, (02) 317 4512, for \$299.

LaserWriter and Windows

I'd like to use an Apple LaserWriter IINT connected to PCs running Windows 3.0 by a serial cable. It's not that I have anything against AppleTalk, but one PC, a Compaq SLT, has no slots, and another has all its slots occupied. We also have visitors with occasional PostScript printing needs who don't want to have to fit a card and install AppleTalk software.

On advice from Apple I had a cable made up —

Printer	PC
2	2
3	3
4	8
5	7
6	4
7	5
8	1
22	9
20	6
DB-25	DE-09

This allows me to use the LaserWriter happily from other DOS environments, including Freelance, WordPerfect, and even Windows 2.0. But from Windows 3.0, I get nothing but a brief flicker from the LaserWriter's communications light.

From Windows 3.0 I've been able to output to a number of other output devices, including an NEC dot matrix parallel printer, a serial HP plotter, a HP LaserJet (on Ethernet), and a parallel TI MicroLaser with PostScript. But, on the LaserWriter, nothing but a brief flicker of the light.

I've checked and re-checked all the communications and printer settings in the Control Panel, and tried several computers, ranging from a venerable IBM AT to a Terran '486 with 5Mb of RAM. I've tried a number of Windows applications from Notepad to Excel 3. I've installed the latest PostScript driver from Microsoft's supplemental Driver Library.

I've sought help from a number of

sources, including Microsoft, and our Apple dealer, but to no avail. They all think it's possible, and can't understand why it doesn't work.

Christopher Deebie

It's a pity that Apple refuses to acknowledge the enormous PC market using its LaserWriters, by not providing a parallel port on their printers. Aside from being much faster than serial connections, the Centronics parallel standard is pretty well adhered-to, unlike the enormous number of variations possible with serial cables.

The trouble with serial cables is that both the computer and printer are classed as 'terminal equipment', and, in order to connect them together, a swap-over (null modem) cable is required. Since this is really a non-standard way of using RS-232 (which was designed for connecting terminals to modems), everybody seems to have their own way of doing it.

The cable diagram you supplied is a pretty standard null modem, for connecting a 25-pin connector at the printer end to a 9-pin AT-style serial port. However, I'd make a couple of changes, which may not improve things, but are useless as they stand at the moment.

The wires connecting pin 22 on the printer to pin 9 on the computer, and pin 8 on the computer to pin 1 on the printer, achieve nothing as they stand, as they are simply connecting two inputs together. I suggest disconnecting both of these, and connecting pins 8 and 20 on the printer end of the cable, and pins 8 and 4 at the computer end. This will ensure that the carrier detect pin is tied in its active state at each end, which probably doesn't matter for printing, but is just playing it safe.

I think the problem really lies in the handshaking. Printers commonly

FEB 1992

use one of two types of handshaking hardware or software. Hardware handshaking involves the printer lowering the signal in one of the special handshaking lines (usually the 'terminal ready' line, or sometimes the 'clear to send' wire), when it is unable to accept any more characters for printing.

Software handshaking uses two special characters sent from the printer to the computer on the return data line, to pause and resume the print stream. This is often called 'X-on/X-off' handshaking.

Windows can be configured to use either system, selected from the Control Panel, in the 'Ports' section. Windows default is no handshaking, which means that it probably won't work at all. The LaserWriter's default is software handshaking, so that is probably worth trying first. Click the X-on/X-off box in the port set-up for COM1, and save it. Make sure the speed is set correctly (default is 9600bps), and also check data format (8 data bits, no parity, one stop bit).

If this doesn't work, then you will need to configure the LaserWriter for hardware handshaking. The following program will do this for you. Enter the lines using any plain ASCII text editor, or use the DOS copy command. Call the file set.pst.

```
statusdict begin 25 sccbatch
exch pop 68 eq {stop} if
serverdict begin 0 exitserver
statusdict begin 25 9600 68
setscbatch
```

To actually configure the printer, enter the following two lines at the DOS command line:

```
mode COM1:9600,n,8,1
copy set.pst COM1
```

These settings are stored in non-volatile memory in the printer, so they should only need to be entered once. If you ever need to change it back to the defaults for

any reason, substitute 64 for 68 in the two places it appears in the above listing, and then send it to the printer as above.

Unfortunately, I don't have ready access to a LaserWriter, but the above program comes straight out of the Apple manual, so it should work. The trick is to get the printer and Windows using the same type of handshaking. Don't forget to enable the RS-232 port on the printer, by setting the DIP switches 1 and 2 down before turning the printer on.

DOS 5 and parallel printers

I recently found that my old faithful printers — an Olivetti Typrinter and a Tandy DMP 400 stopped working the day I installed DOS 5.0 on my machine. Programs in Basic gave me a 'Device Fault' message as soon as they got to an lprint statement — this never used to happen under DOS 4.

When I tried to send output to the printer from the DOS prompt with either Ctrl-P or copy, I got a similar message, which I had observed with DOS 4.01, but it wasn't particularly inconvenient.

The problem affected my other printers, but not the HP Deskjet or other strictly IBM- or Epson-compatible devices. Strange to say, some word processors (such as WordStar), seem still to be able to print normally under DOS 5.0, but not GW Basic or even QBasic, which is bundled with DOS 5.0 in place of GW.

It seems to me that DOS 5.0 uses some of the parallel printer initialisation and handshaking functions which were not so rigorously checked by version 4, and that by disconnecting some of the wires in my printer cable I should be able to prevent DOS from being so particular about the type of printer at the other end.

I have since found that by disconnecting line 31 from the Centronics connector I can get the DMP 400 going.

This is a table of the line functions

according to my printer manuals —

1	Strobe	11	Busy
2	Data bit 0	12	Paper out
3	Data bit 1	13	On line
4	Data bit 2	14	Auto feed
5	Data bit 3	15-18	—
6	Data bit 4	19-30	Ground
7	Data bit 5	31	Reset
8	Data bit 6	32	Error
9	Data bit 7	33-36	—
10	Acknowledge		

Can you please help me figure out how to keep my old printers going? I would prefer not to have to insist that Microsoft give me a refund on an otherwise excellent product.

G. Young

So after saying that Centronics connections were nicely standardised in the previous enquiry, you send me this letter and make a liar out of me!

The fact that the printer behaves differently with different versions of DOS, and works with other software, seems to imply that either DOS 5.0 is forcing one of the outgoing handshaking lines to behave differently, or interprets one of the incoming lines differently to DOS 4.

My guess as to the most likely culprits would be the Busy and Acknowledge lines — the other lines tend to stay in one state or another when the printer is operating, while one or both of these two change state for each character. At least one of these lines is required for proper operation, but with many computer/printer pairings, only one line is needed.

First, try disconnecting the acknowledge line (pin 10 on the Centronics connector), and connect the floating wire to the 5-volt supply on pin 18. If this doesn't help, put the wire back where it came from and attack the busy line. Again, disconnect the wire from the pin (pin 11) and tie this line to ground.

Hopefully one of these solutions will work, although it may not work for the other printers you use — you may need to have a separate cable for this printer.

Unreadable disk

I work as a systems administrator in a government department, and came across a problem that I could not find a solution for. I was hoping you might be able to assist. A document was created in Word for Windows and saved to a floppy disk. All of this was fine. However, when I went to open the file from the floppy disk in the same PC it was created on, the PC came up with a system error when trying to read the disk.

The document will open perfectly from the floppy disk in other PCs, but not in the one it was created on. It is a double-sided high density disk, and a high density drive. When I escaped to DOS to try and 'type' the file, it displayed a message that it could not read a sector of the disk. This seems odd, as the file could be opened on other PCs. I ran Norton's Disk Doctor over the disk, and it found the cluster problem, but no problems with any data files.

Peter O'Farrell

It's not all that uncommon for disks written on one computer to be unreadable on another. Differences in the head alignment of the two drives can mean that the tracks created by one drive, do not lie directly under the heads of another drive. However, it is unusual for a drive not to be able to read a disk which it wrote itself, especially if other drives can read it.

In this case, my guess is that the drive heads are dirty. When the disk was written, the heads were not quite touching the disk surface, so the strength of the magnetic signal on the disk was not all it

should have been. Then, when it was read back, because the heads still weren't touching the disk, they could not recover all of the signal that was there, causing read errors. Presumably, when the disk was inserted in another drive, its heads were cleaner.

At least, that's my guess. I haven't seen anything like this myself, so I'm just speculating. My suggestion would be to clean the drive heads in question, and copy the data from the troublesome disk to another disk, to reduce the risk of it not being readable at some later time.

BBS software

I work for a company that is dedicated to computerising veterinary clinics through a veterinary management system called Netvet, running on Xenix platforms. We are currently planning to set up an electronic bulletin board system as an additional tool in our support area, so users of Netvet systems can leave mail messages for support.

Unfortunately, this company has very little knowledge about the software required for setting up a bulletin board system. If possible, could you please send some information about bulletin board software for both DOS and Xenix?

Cameron Mathers

There is a wide range of bulletin board software around, which ranges widely in complexity. The differences in functionality generally relate to ease of integration into BBS networks, such as Fidonet, and system administration. All provide for multiple message areas and file download areas, and some way to give different access levels and privileges to different users.

We at Your Computer recently were on the lookout for some

new BBS software, to permit our contributors to upload their stories electronically. On the recommendation of Trevor Sheffield from the Sydney PC Users' Group, we finally settled on a system called Simplex, which has everything we need for file and message handling, and is expandable to way beyond what we need. Also, the configuration is menu driven, and user menus can be configured to suit the particular application.

Systems such as Opus, while providing about as much functionality as is humanly possible, are a nightmare to set up for the newcomer, and I doubt that they offer any more useful features than Simplex. Simplex is public domain, and if you can't find it anywhere else, send us \$10 and we'll send it to you on disk.

While there are Xenix-based BBS' around, I have not had any experience with the software, or trying to set up such a system. I suggest you have a look at the latest BBS listing, which we have sent you, for any boards using Xenix as the operating system, and ask the sysops about the software they're using.

If you have a PC problem that's been bugging you, put the details on paper, send them in to Your Computer magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on.

The best Tech Tip published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company. The address in either case is Tech Tips, Your Computer, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

FEB 1992

March 1992

Lost hard disk

I hope that you can help me with my problem — it is my worst to date, and I hope, my last. I have an IBM compatible '386 CPU running at 25MHz, with 1.2Mb and 1.44Mb floppy drives, a 110Mb hard disk, and 4Mb of memory on board. I have been running DOS 5.0 for about a month or so.

For some reason this morning, on trying to start up the computer, I was given the message: Drive not ready, insert DOS disk. Why this has happened is completely beyond me at this time. I had been using the machine for about three hours the day before.

After putting in the DOS 5.0 disk, I was told that the disk is not bootable. I then tried a DOS 4.01 bootable disk, and I did get an acknowledgment with the Microsoft message, and then the A: prompt. And that's where it stops — I cannot get to the C: drive at all — I get a message saying 'incorrect drive specification'.

On rebooting the computer, I pressed DEL to enter the setup program, and according to that, there is no C: drive installed. Being relatively new to the computer game, I am at a loss as to what to do now. I have also tried Spinrite, but it can't find the C: drive either. I live about 30 miles from the nearest town, so I am hoping that this doesn't mean a trip into the 'big smoke' to get the problem solved.

Kevin Manning

Fortunately, your problem is not too difficult to solve, if you know the details about your hard drive. I had the same problem myself a couple of months ago, and it was only when I came across the problem that I realised I knew nothing of the important details of my hard disk.

The hard disk parameters are stored in a special RAM chip in the computer (actually, it's part of the real-time clock chip), which is powered by a small battery when the power is turned off. Sometimes

the battery voltage is a little on the low side (especially after a year or two of use), and the memory won't accurately retain the information which has been stored in it, and the computer forgets the time or date, or forgets that it has a hard disk.

The solution is simple if you know the hard disk parameters — simply enter the setup program (using the DEL key, as you did before), and enter the drive parameters. There are usually 40 or so sets of pre-defined drive types, and if one of these types matches those for your particular drive, simply set the drive type to that drive type. The type numbers and their matching parameters should be listed in your manual, or if not, you can scroll through them on-screen in the setup program.

If there isn't a matching set of drive parameters, then your machine was probably using a 'user-definable' drive type, which most late-model machines have. In this case, you set the drive type to the user-definable setting, usually type 46 or 47, and then manually enter the individual parameters. These are the number of heads, the number of tracks, the number of sectors per track, write-precompensation, and landing zone. You don't need to know the meanings of all of these parameters, only match them to those of your particular drive.

Once you have done this, save the settings using the keys specified in the setup program, and the system will reboot. The hard disk will still have exactly the same data on it as before the computer lost the drive, so you won't need to re-format it.

Incidentally, the reason your DOS 5.0 disk wouldn't boot the machine is probably because it was an upgrade disk, which is non-bootable. Nobody really knows why Microsoft released two variants of DOS 5.0, but it has caused more than its fair share of trouble. Once

your machine is back up and running, use the 'sys' command, as described in the manual, to make a floppy disk bootable, and keep this in a safe place, just in case you ever need it.

Colour printing

I have a problem which I hope you can help me with. About six months ago, I purchased a new Citizen 200GX colour printer. While I have found it an excellent asset, and a very worthwhile purchase, I have had trouble locating a word processor which is capable of taking advantage of its colour printing.

I would like a program which can select different colours, in the same way as I can select underlining. I find it very frustrating to have a colour printer, when to use it as a colour text printer, it must be reset after each line to start a new colour.

Andrew Watts

Although there are a large number of printers on the market, most either emulate, or can be made to emulate, one of the more common makes and models. For example, virtually all laser printers emulate one of the HP LaserJet models — even the new Apple LaserWriters do it. With impact dot-matrix printers, the Epson models are the most commonly emulated.

Nine-pin printers usually emulate the LX-80, or one of the similar models which that company has produced. 24-pin printers are usually compatible with the Epson LQ series. See if your word processor has a driver for one of these printers, and select that as the driver. If there are several drivers, try them all, and see which works best. When doing this, it is a good idea to switch the printer off and on again between drivers, to ensure that the printer is starting from the same point each time.

MAR 1992

However, if the word processor doesn't have any in-built support for colour printing, then you are going to have to define your own printer control strings

Most word processors have the ability of defining 'custom print controls' which give you access to facilities specific to particular printers, which aren't built into the program itself.

Unfortunately, every word processor is different, and you really need to sit down with the manual to figure out how to define and use these custom print controls, and then consult the printer manual to see what codes need to be sent to the printer to enable its special functions. For example, in WordStar, this is done from the wschange program.

Not many word processors have in-built support for colour, although many Windows packages are supporting it. Microsoft Word for Windows, and Ami Pro, both have support for colour text, and provided you have a colour printer driver that works with your printer, you should have no trouble.

Easing the RAM cram

I have an Amstrad PC1286 with 1Mb of RAM. When a program called World Atlas wouldn't run properly, I was informed that the reason had to do with the way that DOS 4.01 apportions the memory. To run the program fully, I was told I needed an earlier version of DOS that used up less RAM.

Another program — Publish It Lite, when loaded, leaves only 44Kb spare. Since two images on a page (in this program) can use 30Kb, then it doesn't take long to run out of memory.

Why is it, that with 1Mb of RAM, the memory is running out, with programs that need 640Kb or less. Can it all be blamed on DOS 4.01, or is there some other reason?

J. Delamoir

Tech Tip of the Month Remote install in LapLink III

I RUN A SMALL computer consultancy and service business, where sometimes the odd or unusual problem arises. So far I have managed to solve them all, (touch wood). The most common problem is users playing with their autoexec.bat or config.sys files and ending up at the C: prompt wondering what to do next. Anyway, the reason for this letter is a problem that might interest you.

The problem involved an Epson PCe XT, in which the floppy disk controller went into retirement, and, being part of the motherboard, that meant quite an expensive replacement. As the owner was buying a new computer, this would solve the problem. Ha! On the XT were all his program files for his business, hundreds of data files, plus numerous other files, except something like ll.exe or ll3.exe. Not having a modem program ruled out transferring through modems, which would have been one way.

I diverge here to mention that backing-up was not a regular practice with this customer, but certainly is now. LapLink III was the next alternative, which as you know is supposed to transfer the ll3.exe to a computer that doesn't have it on its disk already. As I said, it is supposed to. Following the instructions explicitly (and not so explicitly), no matter what I did, no file transfer took place. Many phone calls later I was no better off, and wondering how some experts get to be called experts.

I then hit upon the idea to maybe create a batch file to push the file over (if push is the right word). Lo and behold one ll3.exe file transferred — the relief was overwhelming.

All I did was make a batch file containing the commands that are supposed to be typed on the PC that is to receive the ll3.exe file, and copy the file to com1.

```
rem GO.BAT

mode com1:2400,n8,1,p

city com1
```

I hope this information may be of some use to somebody in some way. Keep up the good work with *Your Computer* — I can honestly say I read it from cover to cover. It is not only interesting, but informative and comprehensive. Well done!

Kevin J. Kershaw

We use LapLink quite frequently here in the *Your Computer* offices, to transfer our benchmarks and favourite applications over to test machines, and have often encountered the same problem with the auto-install facility. I must admit that using a batch file never occurred to us, but we'll remember it in the future.

By the way, I assume you meant 'run the batch file' in the second-last paragraph, not copying it to the COM port. And thanks for sending the tip in on disk — it certainly made my job that little bit easier.

The ultimate culprit for the 640Kb RAM limit which we're all burdened with at the moment is myopia on the part of the designers of the PC's original processor chip — the 8088. This chip was designed really as a souped-up 8080, an old 8-bit processor, with a memory

limit of 64Kb. The designers added segment registers and a few other bits, and expanded the memory addressing capacity of the chip to 1Mb.

This decision (together with the decisions of IBM to use it in their PC, and of Microsoft to base their

MAR 1992

operating system on it), have left us with two legacies. The first is the segmented memory structure, which is a pig for programmers to work with, as the memory is still addressed in 64Kb blocks. While this segment limit is invisible to end-users, it wastes programmers' time which could be better put toward other things.

The other problem is that it left no way open to expand to bigger memory systems in the future. The 80286 can address up to 16Mb of RAM, but in order to do so, it has to be switched into its so-called 'protected mode', which is incompatible with DOS programs. So this extra memory is only usable for storing data, not as program memory.

So, while 640Kb of RAM seemed like a lot at the time the PC came out (the rest of the 1Mb addressing range of the chip is taken up by system ROM, video adapters, and other peripherals), software quickly grew to fill the space available. As DOS grew in functionality (and size), programs which would fit comfortably in memory with previous versions of DOS, suddenly started getting squeezed for space.

Then, along came DOS 5.0, which reversed the trend by using a little-used bit of memory that '286 and '386 chips can utilise, while still being compatible with DOS programs. So, with DOS 5.0, there is actually more conventional memory available to programs than at any earlier time in DOS' history — over 600Kb in some cases.

So instead of going back to an older version of DOS, I'd suggest upgrading to DOS 5.0, and gaining use of at least part of that extra 384Kb of RAM that you have. And you can still use the rest of the extended memory for other things, like a disk cache or a small RAM disk.

Starting a BBS

For the past 18 months, I have been a regular member of the various bulletin boards that are located in my local area (the NSW Central Coast). Having used the BBS' available, I am now interested in beginning my own BBS. It would be appreciated if you could recommend any books or literature that is available on beginning bulletin boards. I am particularly interested in discussions on programming, file compression, doorways (for on-line games), and user groups.

Keith Billington

Since you are already active on a number of your local boards, the best place to start is probably by asking some of the local sysops what they'd recommend. Sysops tend to have their own favourite BBS software (the package they use on their own system, of course), but most will be honest enough to tell you how easy it will be to get up and running.

In addition, if you use the same software that other local people are using, then you have a ready pool of knowledge to draw on when the inevitable questions arise during the setting-up stages, or if you want to change your configuration later on. Sysops generally don't mind helping you in this way, since once you become familiar with the software, they might be able to call on your help one day.

The financial aspects of running a bulletin board should not be overlooked either. Aside from the initial capital costs of the equipment required, there are the on-going costs of a dedicated phone line, and the occasional unforeseen disaster (hard disk failure is a common killer of BBS operations).

Many sysops charge dues for membership to try and cover these costs, but rarely do they come

anywhere near covering the running costs of the BBS. But, if you're prepared to invest the time and money, it can be a rewarding (if not frustrating), pastime.

Graphic file formats

In your November 1991 issue, you ran a piece entitled 'Standard forms with Ventura'. My attention was caught by the words '... saved each graph as a pic file ...'. My question is: what is a pic file? Is it just another name for a Mac file for PCs? Is there a standard pic file format?

I ask this because I recently came across a pic file and was unable to view it. I purchased a book called Bit-mapped Graphics by Steve Rimmer, and although it explained what Mac, gif and pcx files are, it did not have any information on pic files, and how one would go about decoding or unpacking the picture. I have tried to read the header using graphic conversion programs, but to no avail.

I believe that this particular file was created by a package called IBM Storyboard. If you could give me any information on this file format, or point me towards a reference book, I would be most grateful.

Patrick Clift

The number of graphic file formats seems to be without end, and while a few formats are establishing themselves as standards (GIF, TIFF, and PCX, seem to be the most common on PCs), the situation is not helped when several different file formats use the same name and file extension.

I can think of four programs which use pic files, and all have their own file format. Lotus uses the pic extension to save graphics files for printing with its PrintGraph utility. Mouse Systems has (or at least, had) a program called PC Paint, which used its own pic

MAR 1992

format, as does IBM's Storyboard. And then there's the Mac's pic file format, which is different again.

Ventura can read the PC Paint format pic file, which is what the article you mention was referring to. However, I think only Storyboard is capable of reading Storyboard's pic format. Even its pcx file format seems to be different to everybody else's.

The only thing I can suggest is to get a copy of Storyboard, load the file, and then save it in some other format that something else can read, such as GIF.

BIOS drive support

I have recently purchased a '386 motherboard with an old Phoenix BIOS on it. The problem is that it does not have support for my hard disk in the CMOS setup, and there is no user-definable hard disk setup. What do I do? After spending half a day ringing around, I have been told that I can't just get the BIOS. Would anyone at Your Computer know from where I might be able to purchase an updated BIOS for my '386 machine?

Mark Seifein

I am not aware of anybody currently selling updated BIOSes in Australia, but there is a company in the United States which does. It is called 'Upgrades Etc.' and its address is 15251 NE 90th Street, Redmond, Washington, 98052, and its fax number is +1 206 881 8294. The only phone number I have for it is a US 800 number, which is not accessible from Australia.

This company sells licensed BIOSes from Phoenix, Award, and AMI, for a wide range of motherboards, at prices ranging from US\$60 to US\$80. I asked if it had an Australian distributor, but it doesn't, (at the moment).

If you possess an EPROM programmer, or have ready access to one, it is possible to alter the built-in drive table, changing one of the existing drive types to match your own drive. The location and format of the drive table is pretty easy to find out in most BIOSes, but there's not really the space to go into it here. The only possible problem with this approach is that the timing of early AMI BIOSes was not compatible with some IDE drives, but as you've got a Phoenix, this shouldn't be a problem.

Printing graphics characters

In your November issue was an article entitled 'Teach your printer to fly', and there, I thought, was going to be my hope and salvation, as far as printers are concerned. Most of the material I had read and heard of in fragments from here, there, and everywhere, but not in such a collated and concise form.

So I had another go at printer controls, but alas, the result is still 'no go', and I don't understand why. I have an Epson RX-100 printer, and try as I might, the ASCII character set from 127 to 255 eludes printing. Sure, I can get the characters on the screen, but not printing.

The system is an IBM-alike, running at 4.77MHz, with a 30Mb hard disk. The autoexec.bat file contains the ansi.sys driver, also.

Frank Damen

First of all, I must point out that the characters above 127 decimal are not strictly ASCII. The ASCII character set only defines those characters which can be represented by a seven-bit binary number — that is, those from 0 to 127. Many computer manufacturers, since their machines deal with information in multiples of eight bits, extended the character set to 256 characters, and of course,

IBM's extended character set has become the standard 'extended-ASCII' character set, but is by no means the only one in existence.

While this distinction might appear merely one of semantics, it does explain why some printers do not support the extended IBM character set, either ignoring the eighth data bit entirely, or printing a completely different character set. However, most late model printers do support the IBM character set, although in many cases, you have to set an internal DIP switch to enable this feature.

I'm not familiar with the particular Epson model of which you speak, but if you refer to the section in the manual which deals with DIP switch settings, you will see if there is any switch to allow you to enable this feature. If so, turn it on — it shouldn't affect printer operation in any way, except to allow printing of the extended characters when called on so to do.

Failing that, there is really nothing that you can do. I have an old Epson LX-80 which doesn't support the extended character set, so you're not alone. By the way, ansi.sys has no effect on the printer — it is an enhanced screen and keyboard driver. Also, it should be in config.sys as a 'device=' statement, not in autoexec.bat.

Enhancing DTP printing

I am writing to you for some information concerning a statement made in the article 'Printing for DTP' in your March 1991 issue. The statement is made that a package such as Publish It! running on an XT outputting to a good dot-matrix printer with print enhancement software will give satisfactory results.

I have an Amstrad 2086, and an Epson LQ500 running Publish It! I live in the country, in an area poorly supplied

with literate computer suppliers, which forces me to have to shop by mail or phone. Whereas in the city you can make 10 phone calls for \$3, I can spend that much just listening to music on hold.

Could you recommend a print enhancement package that would suit someone like me?

Merwyn Passmore

One popular DTP print enhancer is Publishers' Powerpack, which gives scalable fonts to users of non-PostScript printers.

I don't know how well it works with Publish It!, but the distributors, Trio Technology, on (03) 585 0566, should be able to tell you without keeping you on the phone for too long.

Another possibility lies in a program called GoScript, which takes a PostScript file that has been printed to disk, and converts it to a bit map, for printing on printers such as the LQ series.

This will work with any program that can print a PostScript file to disk. It, too, is distributed by Trio Technology.

April 1992

Not enough memory

I have a '486/33 with 640Kb of base memory, plus 4Mb of extended RAM, and I still get messages that say 'not enough memory'. What I don't know is how to utilise this extra 4Mb so that I don't get this message and have enough memory to run the various programs.

I have enclosed copies of my auto-exec.bat and config.sys files, which may help determine where I am going wrong.

Also, if I load my mouse driver to use in a program, and then try to load Windows after I have finished in the previous program, Windows locks up and will not run until I have gotten rid of the mouse driver previously loaded. Your help would be most appreciated, and hopefully stop my frustration.

Peter Greenwood

First of all, the 4Mb figure is almost certainly the total amount of memory in the machine, not just the amount of extended memory. Of this total, 640Kb is the base memory, and 384Kb is probably reserved for shadow memory use, (copying code from slow ROM chips to RAM where it will execute faster), leaving 3Mb of extended memory.

From your config.sys file, you have 1Mb of this allocated to a RAM drive, and 2Mb allocated to a disk cache, leaving none for use by

any other applications. Windows can reduce the size of the disk cache if it needs to, to gain use of up to 1.5Mb of extended memory. This is probably not the best utilisation of available memory.

For starters, a RAM drive is usually not a very efficient use of memory. Unless you have applications which make extensive use of temporary files (such as most Windows applications), then the speed advantage of a RAM drive is likely to be minimal. Also, make sure that these temporary files are being redirected to the RAM drive, or you won't get any speed improvement at all. For Windows applications, this is achieved by adding the line:

```
SET TEMP=D:\
```

to your autoexec.bat file, where D: is the letter of your RAM drive.

If you do decide to keep the RAM drive, determine the amount of space that you really need, and reduce the size of the RAM drive to this figure (plus a little bit for unforeseen circumstances), otherwise it's just wasting space.

A disk cache, on the other hand, requires no manual intervention in order to make use of it — it will keep the most-accessed areas of the disk in memory, ready to be accessed instantly. Again, I think 2Mb is a bit of an overkill — I'd use only half that amount.

One further thing that will give you a few Kb of extra memory is to load all your device drivers 'high', by changing all the 'device' statements in your config.sys file, to 'devicehigh'. This will load the drivers into vacant areas between 640Kb and 1Mb, which were made available with the 'DOS=high,umb' line in your config.sys file.

These measures will maximise the amount of conventional memory available to all applications, and increase the amount of free extended (XMS) memory. Most applications can't use extended memory, so it won't help them.

If your applications can use expanded (EMS) memory, then adding the line:

```
DEVICEHIGH=C:\DOS\EMM386.EXE
```

to config.sys will allow unused extended memory to be used as EMS by applications which need it.

Your mouse problem is a result of Windows trying to be an operating system while being built on top of another operating system, (DOS). The logical interface to any hardware device should be part of the operating system, (either in-built, or as an add-in extension, installed with a 'device=' line). Microsoft has taken it upon itself to implement mouse drivers from within Windows, which often causes conflicts with drivers already installed under DOS.

APR 1992

The easiest solution for most people is to run programs that need the mouse driver from batch files, and then use two little utilities called Mark and Release to allow the mouse driver to be removed from memory when it is not needed. Mark leaves a little 'marker' in memory before any TSR (terminate and stay resident), programs (like your mouse driver), are loaded. When Release is called, it looks for the marker, and then removes everything installed from there on, freeing up all the memory taken up by those programs. They also store and retrieve a copy of the system's interrupt vector table, so that programs that hook themselves into various interrupts, like mouse drivers, can be removed without leaving the rest of the operating system in limbo.

Mark and Release are part of a handy little utility package of copyrighted, but free, software, available from many bulletin boards as TSRCOMxxZIP where xx is the version number, (currently 3.2). If you don't have a modem, send us \$10 for disk duplication, and we'll send you a copy, (specify what you want, of course).

The archive also contains a number of other utility programs, the most useful of which is *mapmem*, a little thing which details which programs are using your memory.

BIOS upgrade needed

Two years ago I purchased a fairly basic '386 system from a local supplier, and I have since increased the machine's capacity in many ways; in fact nearly everything has been replaced. I have upgraded the hard drive from 40- to 200Mb, installed a Tseng ET4000 VGA card, and added a few other things too.

I have over this last two years noticed a number of inconsistencies with the

system (too numerous to list here), and they all appear to point the finger of doom at the BIOS being a tad incompatible. So, having spent a small fortune on various improvements, I have decided to spend more and upgrade the BIOS, (yes, I know it would be easier to just sell the dog, but I can't bring myself to pass my problems on to some other unsuspecting person). This is where I run into trouble and is why I resorted to contacting your (much read), magazine.

My present BIOS is labelled DTK by Datatech Enterprises, Version 4.26 and the Board Chipset is labelled PEM; both of these appear to be virtually unknown in this country. Hence my question — 'Where can I obtain a upgrade?' or alternatively, 'Do you have a contact address for Datatech Enterprises?'

Mark R Standing

There is a company in the US which specialises in BIOS upgrades for old machines, which I have mentioned before in 'Tech Tips'.

Its range includes DTK chips, but last time I checked it did not have an outlet in Australia, nor am I aware of any similar company operating here.

The company is called Upgrades Etc, at 15251 NE 90th Street, Redmond, WA, 98052 USA; fax +1 206 881 8294.

Executable Basic files

I write programs in GWBasic, and believe that they are good enough for shareware distribution. I really would like to distribute them, except that they cannot be loaded from DOS. At the moment I get around this by using a batch file.

Is there a way of converting these files into executable files, or is there a way of loading them without going through GWBasic?

Ryan M. Shulman

GWBasic is a Basic interpreter — that is, in order to run a program, the interpreter (gwbasic.exe), reads the program line by line, and then acts on the instructions. This happens each time the program is run, so the interpreter needs to be available whenever you want to run the program.

In order to convert your Basic programs to machine-readable executable files, you need a Basic compiler. Unlike an interpreter, a compiler performs a once-only operation, converting the basic program into machine-code instructions that the processor can execute directly. Once a program has been compiled, it becomes a stand-alone entity, which can be run on any (compatible), DOS machine.

Two common (and inexpensive), Basic compilers on the market at the moment are Borland's Turbo Basic, and Microsoft's Quick Basic. Both these compilers are much more sophisticated than GWBasic, and your programs will probably need some modifications in order to compile correctly, but the programs will then be stand-alone exe files, and run much faster than they did under GWBasic.

Note: the QuickBasic shipping with 5.0 is not the compiler, but an interpreter using the QuickBasic language.

Programs written to run under this interpreter should compile under the QuickBasic compiler with little or no modification, but you'll still have to buy the compiler as a separate item.

Coherent Users' Group

In the May 1991 issue, we printed a letter from a reader asking if there was a users' group for the Coherent operating system. At the time we were unaware of any such group,

APR 1992

but we've just received word that such a group has just been formed in Australia.

The Australian Coherent User group (ACUG), has been founded with the co-operation of the authors of the Coherent Operating System, the Mark Williams Company in Chicago, USA. The intent behind the establishment of ACUG is to provide a forum where information exchange regarding the use of and programming under Coherent can take place, without the difficulty and expense of directly contacting the MWC BBS in America.

In order to facilitate the distribution of both Coherent updates, and public domain Coherent software, MWC is providing a monthly distribution of the contents of the MWC BBS on disk to ACUG, thereby maintaining a mirror of the MWC BBS downloads directories for the benefit of Australian Coherent users. These files are available by direct UUCP transfer from the ACUG BBS. Registered Coherent users are invited to join ACUG, and to access this information.

The ACUG BBS will also provide email exchange between ACUG members, as well as the posting of the comp.os.coherent newsgroup to those with news software installed on their computer. Furthermore, local ACUG newsgroups will be available for use by members. These groups will include (but not be limited to), binaries, bug reports, source code, and update information.

Membership of ACUG is available to all legitimate Coherent users, subject to the submission to ACUG of a unique and valid Coherent Serial Number. The serial number will be used as an individual's password for accessing the ACUG BBS. Following discussions with the administrators in Chicago, it has been agreed that an annual subscription of \$30 per year

be set to offset some of the overheads associated with running the ACUG BBS in Melbourne.

Anyone interested in joining and availing themselves of the facilities offered should email acug@spectre.pub.uu.oz.au, requesting information.

Alternatively, log onto spectre (the ACUG BBS on (03) 882 1558 at 2400bps), using UUCP as 'uguest' (no password necessary), and forward email to acug@spectre, again requesting information. Please ensure that you leave appropriate contact details. As a last resort, telephone Richard Lindner on (03) 824 0655, or (03) 882 4690 after hours, or fax him on (03) 827 4209.

Coherent is a multitasking multi-user operating system written and marketed by the Mark Williams Company in Chicago. It is modelled after the Unix operating system, and provides almost all the functionality and utilities of Unix. The porting of most utilities from one of the implementations of Unix to Coherent is typically a trivial task.

Accessing XT memory

I have an XT clone with a 384Kb RAM disk set up above 640Kb, and have some questions as to how the memory is organised, and how to use debug above 640Kb.

Is the RAM from 640 to 1024Kb continuous? Does it have gaps in it for the hard disk ROM and system ROM, or is there RAM in parallel with the ROM?

Is the adapter segment RAM (640 to 1024Kb), the same area as used for the RAM disk in my machine?

Also, debug seems to be able to read memory from 640 to 1024Kb, but is unable to write to the sections of memory which are RAM. Is this correct; if so, why?

John Badger

Not many clones have more than 1Mb of memory, because the processor (the 8088), can only address a total of 1Mb of memory, and that includes system RAM, ROM, and various peripheral cards such as video adapters. IBM filled the memory up to the 640Kb boundary with RAM, leaving the rest of the memory space free for the other devices. So if you have 1Mb of memory, the extra 384Kb has to go somewhere else.

AT-class machines do this by locating the extra memory above the 1Mb barrier. But there is no such thing in an XT — the memory space stops at 1Mb. The most common way to address extra memory in machines such as this is to use bank-switching, to split the extra memory into small chunks, and switch these chunks (called pages), into a vacant area of the memory map between 640Kb and 1Mb, (there are large gaps of unused memory space in this region). Not all the memory can be accessed at once, but for storing data (or disk information), it is quite acceptable.

The Lotus-Intel-Microsoft Expanded Memory Specification (LIM EMS), is the most common bank-switching scheme in use, and is likely the one used in your machine. This divides the extra memory up into 16Kb pages, any four of which can be switched into a 64Kb page frame in a vacant part of the memory map between the top of conventional RAM and the end of memory — E000 to EFFF (hex) is a common place, just below the system ROM.

Other than the page frame, all of the other regions in this area of memory are either ROM, video RAM (the location and size of which varies with the type of adapter installed), or unused. You can't use debug (or anything else), to write to the ROM or unused memory areas, but you should be able to write directly into the video RAM or the EMS page frame.

APR 1992

However, don't forget that the video RAM is also being written to by DOS, so any values which you write there might be overwritten immediately by DOS, as it updates the screen. Also, writing to the page frame is pretty hit-and-miss, since you won't know which pages of memory you are writing to, unless you check the page selection registers first, (and they don't change after you read them).

Full SuperStor disk

After reading the review of SuperStor in the October issue last year, I decided it was just what my old 40Mb hard disk needed. I dutifully obtained a copy, and installed it on my hard disk, and it has performed flawlessly, until now. I was trying to save a file from Excel, when I was given a message that the disk was full. Fair enough, I'd furiously been putting stuff back on it that I had archived on floppy because of lack of space before.

However, I soon discovered that it had left a lot of bad clusters on the disk (as reported by chkdsk), and there was no free space available, even though it hadn't saved my spreadsheet, or even part of it. How can my hard disk suddenly develop bad clusters like this? The SuperStor manual says that they will go away when I reboot, but they haven't, and I still can't put anything else on the drive.

R Evans

The so-called bad clusters are a result of SuperStor's operation, but they are not really bad, and you can recover from this situation without too much trouble. What SuperStor has done is mark these sectors as bad, because it over-estimated the amount of space left on the drive, which has now filled up before SuperStor expected it to.

At first, this might seem like a rather kludgy way to tell DOS that

the disk is full, but really, there isn't a better way for SuperStor to tell the operating system. When SuperStor reports the amount of free disk space to DOS, it is only an estimated amount, the true amount of free space will depend on how compressible the data is that is going to be stored there.

So, SuperStor (and therefore, DOS), doesn't know exactly when the disk will fill up, so it could conceivably fill up when DOS thinks there is still space available. For example, if SuperStor reports that there is 4Kb of space left on the drive, and tries to write 4Kb of uncompressible data to the disk, there won't be room, since SuperStor bases its estimate on, approximately, a 2:1 compression ratio.

But once SuperStor has reported this amount of space available, it can't go back on its word and say that there is only really 2Kb free, so (apparently), the only way out is to mark those sectors of the disk as bad, so DOS will get a write error.

SuperStor actually comes with two utilities — Sshrink and Ssxpand, which reduce or increase the size of the file containing the virtual compressed disk. The documentation fails to mention these two utilities, and the fact that they can be used to fix a partition after SuperStor has marked the clusters as bad. Even worse, the manual is incorrect in saying that this problem is automatically fixed when the system is rebooted. You need to use one of these utilities to repair the disk.

If your SuperStor drive is drive D: (for example), then type 'dismount D:', to make the drive appear as a normal device again, and then type:

```
SSHRIK D:\SSPARTSS.ADD 1
```

which will shrink the size of your SuperStor partition by 1Kb, but more importantly, unmark the so-

called bad blocks. You can then run Ssxpand (with the same command-line parameters), to expand the file back out to its original size. Then, remount the drive, or just reboot.

The writers of SuperStor seem to have taken the same attitude to users as IBM has taken in relation to low-level formatting hard disks — it's too technical for the user, let's leave it for the experts, (who can then charge exorbitant sums of money for their trouble). There are enough tricks to operating a computer, without any additional ones created by deliberately withholding information such as this.

Commodore RAM cram

I'm writing to you in the hope that you will be able to help me with a few problems I'm having with our Commodore PC 40 AT-compatible with a 40Mb hard disk, 1.2Mb and 720Kb floppy drives, and 1Mb of RAM, (arranged as 640Kb of conventional and 384Kb of extended). The DOS is version 3.3.

The main problems are with access to the extended memory and TSR programs. When loading some of our simulator games, a message appears on the screen that there isn't enough memory to load all the graphics. I went into Norton Utilities' System Info version 5, and checked that the machine was recognising the 1024Kb of memory, which it was. I then checked the memory usage, and found that it was really only working with 640Kb, 82Kb of which was used by TSRs, leaving only 558Kb to play with.

At the rear of the machine there is a switch which enables and disables the extended memory, and it is in the 'up' position, which I presume means that the extended memory is enabled. How do I gain access to, or 'get at', the extended memory?

I am also presuming that if I can use the full 1024Kb of memory available, then 82Kb of TSRs won't really bother the loading of graphics or other programs. Is

APR 1992

this correct? Does it matter which drive I use? — all the games are on drive D: at the moment.

I tried to include himem.sys, but all that managed to do was freeze the system immediately after bootup, so I deleted it. Unfortunately, being on Thursday Island (32km from the mainland), there is no access to computer stores, or people with expertise. I hope that you can understand all this and are able to assist me.

Anne Zinnegger

The first thing that you should do is to upgrade your DOS version to 5.0 — this version has the ability to use up to 64Kb of your extended memory to give more of that 'bottom 640Kb' up for use by your TSRs and other applications. Unfortunately, there is no way that most applications are able to use the remainder of extended memory, with few exceptions, such as RAM disks or a disk cache.

If your system has a NEAT chipset (New Enhanced AT, a set of integrated circuits which greatly enhances the '286 processor), then you could convert some of this remaining memory into 'high RAM' and load some or all of your TSRs into this area. If your computer's power-on message or documentation mentions 'NEAT', then you're in luck.

Failing that, if you still don't have enough room for both your TSRs and games, then you will need to remove the TSRs from memory before running the games. One way to do this is to use the Mark and Release utilities mentioned earlier in this month's column. Another possibility is a shareware program, Boot.sys, which allows you to select different options, (such as TSRs each time the machine is re-booted). This is available from many BBS', or from its distributor — Budgetware, on (02) 519 4233. There's a review of Boot.sys in the November 1991 edition of 'Tech Tips'.

Low-level formatting for IBMs

In your January 1992 issue you advised your correspondent how to do a low-level format on his XT clone, using debug, and typing 'g=C800:5'. I have a similar problem with a 10Mb hard disk. I cannot access it, yet with Checkit, I am told that cylinder 0 is marked with a low-level format, cylinder 1 is too damaged to report, as are cylinders 8 and 240.

Is there a similar command for debug with the IX (DOS 3.2) to low-level format the hard disk? Neither the hard disk manager or fdisk will access it. Checkit also reports that it's a non-DOS disk.

Barry Irwin

The debug command 'g=C800:5' is really a 'goto' command, telling debug to start executing the program located at that address. Almost all clone hard disk controllers have a low-level format program at that address, but alas, the genuine IBM models do not. IBM seems to

think that low-level formatting is not the sort of thing an end-user should do, and puts its low-level formatter on the Advanced Diagnostics disk for that machine.

Fortunately, third-party programs have low-level formatters in their suites. Best known are Disk Manager, once provided with many hard disks in order to allow DOS 3.2 to access extended partitions on the drive, and Landmark's Probe.

Once the disk is formatted, run fdisk, and set up a primary DOS partition. If there is still a problem with track 0, this will show up at this point, if it hasn't already. If cylinder 0 is damaged, there's not much you can do to save the drive. However, if it is cylinder 1, you can create a dummy partition of one cylinder, and then put your DOS partition after this one, and select the second partition as the boot partition. This was detailed in the May 1991 issue.

Probe comes from Interworld Electronics, (03) 563 7066, and Disk Manager from Software Express, (03) 663 6580.

APR 1992

Tech Tip of the Month — Keyboard assignments

IF, LIKE ME, you are constantly using DOS to do such mundane tasks as finding and copying files, then you are probably sick to death of typing 'dir /w' or 'dir /p'. If so, I suggest assigning the strings to an easily-typed key, as demonstrated in YC, August 1990. But, on extended keyboards, these keys are inconvenient, and hard to type quickly.

I decided to assign the command to Control-Enter, which is very easy to type quickly, and very simple to program. First, you need to load ansi.sys into memory. This is performed by adding the line:

```
DEVICE=C:\DOS\ANSI.SYS
```

to your config.sys file. If you are using DOS 5.0, and have DOS loaded high, then use a 'devicehigh' statement instead. Then add the following to autoexec.bat:

```
ECHO<esc>[10;"DIR/W";13p
```

Note that <esc> means the escape

character. To enter this character using the DOS editor (or WordStar), type Ctrl-P before pressing the escape key. In Microsoft Word, hold down the Alt key while typing '27' on the numeric keypad, (you must edit the file in format).

After making these changes, reboot the computer, and whenever you press Ctrl-Enter at the command prompt, the computer will act as if you had typed 'dir /w'. By changing the string between the quote marks, any command can be assigned to this key combination. It may take a while to get yourself out of the habit of typing the string manually, but you will save time eventually.

Cameron Walker

This technique can be used to assign virtually any key combination on the keyboard. A full list of the key codes (these replace the '10' in the above statement), is in chapter 15 of the DOS 5.0 manual, including the F11 and F12 keys, and the other extra keys on extended keyboards.

May 1992

Turbo speed isn't

I have an older model IBM XT that has a 20Mb hard drive (made in 1978, but still going strong), a CGA monitor, and two 360Kb floppy drives. Because of a problem with the original motherboard, it was replaced with a GT-20 Mini 10MHz motherboard, which has a Modular BIOS, ver 3.1jk, from Award software, and is fitted with 640Kb of memory.

When the unit is booted up with the turbo switch off, it performs normally, and does not give any problems. In this mode, it is very slow of course, and makes some programs impossible to use effectively.

If the unit is booted up with the turbo switch on, it does not run properly at all. It appears that the autoexec.bat file is not completely executed, and the system cannot always find a specified directory. If a directory is found and the program runs, upon exiting, the system cannot find command.com, or indicates that it is invalid and the system is halted.

I have set the comspec variable in auto-exec.bat, reverted to DOS 3.3, and copied the command.com file to every conceivable directory, all to no avail. If the system is run without the turbo switch on, the program will operate normally, but when it is exited the problem appears.

I have checked all jumper and switch configurations on the motherboard and cards, and can find no conflicts. I would be pleased to receive any advice that you may consider relevant.

J.B. Price

Your problem sounds like a case of one or more of the expansion boards in the computer not being able to run as fast as the new motherboard. The designers of the add-on cards, such as the hard and floppy disk controllers, designed them to work at the speed of the original motherboard. When called upon to operate at the higher speed, it only takes one component to throw up its hands in disgust, and it's not going to work. The only real

way to check this is to swap the boards for ones which are known to work at the higher speed.

Another possibility is the RAM on the board, especially if you transferred it from the original XT motherboard to the new one. Like expansion cards, the RAM chips have to be able to work as fast as the processor, otherwise data errors will occur when they are being read. This usually results in parity errors, but some motherboards don't use parity checking on the RAM, so it could go unnoticed until it corrupts some critical instruction and falls over.

For a 10MHz XT motherboard, 120ns (nanosecond) chips are fast enough — indicated by a '-12' suffix after the chip's number. If the chips are slower than this (indicated by a higher suffix), then replace them with faster (12ns), chips.

Unfortunately, other than substitution, there is no way to check which (if any), of the expansion boards is causing the trouble, but it is not an uncommon problem. If you have a friend with a turbo XT (that works), you could try swapping boards between them, to see if you can find which board is too slow, and then replace it with one that is guaranteed to work at the faster speed.

The 'missing command.com' message sounds like a separate problem, especially since it happens at both speeds. Make sure that all copies of command.com came from the same DOS version that the hard disk is formatted with, and double-check that the comspec environment variable is set correctly. A common mistake is omitting the drive letter, and if you terminate a program on a drive other than the one on which command.com resides, DOS won't be able to find it.

As a last resort, try increasing the number of file handles in config.sys, in case the application is not closing its files properly, leaving DOS no free file handles with which to re-load command.com.

Disk problems and solutions

I'm running an ADV '286 with 1Mb of RAM and an AMI BIOS. It has a 1.2Mb A: drive, a 1.44Mb B: drive, and a 42Mb hard disk, as drives C: and D:. Drive B: is the problem — when I do a directory, there seems to be no serial number. Is this a quirk of DOS 4.01, or just my drive?

The second problem is that I can't get the computer to recognise a disk in it on start up, and boot from it. Drive A: can do this, why can't drive B:?

*My next question is about the Chkdsk command — if I'm in the DOS directory and I run Chkdsk with the parameter *, at the end I get a message that looks like this:*

```
C:\DOS\LONK.EXE Contains 5 non-
contiguous blocks
C:\DOS\SYS.COM Contains 6 non-
contiguous blocks
```

(the actual listing is much longer than this). What does this mean? If I'm in a directory that has no non-contiguous blocks, I get the message:

All specified file(s) are contiguous

What does this mean?

Last question — whatever happened to IBM's PVGA graphics card? What was wrong with it, and was it released here?

Here's a quick tip for people who use the Subst command. In DOS there is an undocumented command — Truename, (I say undocumented, as there is no mention of it in the DOS 4.01 manual). You can use this command to find out what the Substituted drive was originally. For example, if you used Subst to make a certain directory appear as drive E:, then typing Truename from the E: drive (or typing Truename E:), will return the original path name.

Stuart George

First things first, I wouldn't worry about not getting the serial number from the B: drive — I have yet to

MAY 1992

see any application for it, aside from filling up another line on the screen. My guess is that you're using pre-formatted disks in the B: drive — all the pre-formatted ones I've seen don't return a serial number (both 3.5- and 5.25-inch sizes). I don't know why this is, but it's probably related to the way in which the disks are mass formatted during manufacture. If you really want a serial number, re-format them all, but I wouldn't bother.

The reason that your computer won't boot from the B: drive is that the BIOS doesn't try to. This is true of most BIOSes, and is probably a carry-over from the days when people had two identical drives, so there was normally no reason to boot from the second drive. With the advent of 3.5-inch drives, this would be a worthwhile addition to any BIOS, instead of having to swap the drives over if you have a 3.5-inch disk that you absolutely have to boot from, like I did when installing OS/2 2.0.

The only exception to this rule that I've seen is a Philips AT, which allowed booting from either floppy drive, and if you booted from the second drive, the BIOS would swap the logical drive assignments over, so that the drive you booted from would always be drive A:.

Contiguous files are files where all the clusters that make up the files on the disk immediately follow the previous cluster. While this is the most sensible way for a file to be, in order to be accessed quickly, it's rarely that a large file can be written to a disk in a contiguous manner.

When DOS writes a file to the disk, it looks for the first available cluster, and puts the first cluster of the file there. It then looks for the next free cluster on the disk, to put the next cluster of the file. If this cluster follows the previous one, then the clusters will be contiguous, but if the cluster is somewhere else, then the file will be fragmented at that point.

While DOS has a built-in mechanism to detect fragmented files, there is no way to de-fragment the files, without backing up and restoring the whole drive, or using an external third-party product, such as Norton's Speed Disk (part of the Norton Utilities suite). Fragmented files slow down disk accesses, by making the drive heads move back and forth a lot more than they would otherwise need to, in order to read the various clusters.

I don't know about IBM's PVGA — you're not thinking about the Professional Graphics Controller (PGC), by any chance? That card was probably ahead of its time, as far as the average PC user was concerned. It was the first card from IBM to have an on-board processor, but was probably faster (and more expensive), than most people needed at the time, before PC applications had much in the way of graphics work. If there was another card called the PVGA, it appears to have suffered a similar fate.

And thanks for the tip about 'Truename' — DOS is notoriously poorly documented, but actually leaving out commands is going a bit far. Incidentally, the command also works with drives that have been grafted to the directory tree using the Join command.

Undeletable directory

I have installed an accounting program in a directory called 'select', containing 18 files. As I do not like the program, I decided to delete the files and remove the directory as well.

Unfortunately, I can't do it. The shell's directory tree shows the 'select' directory, but the file section says 'no files in selected directory'. I tried the search option and the editor to try and find the

files, but to no avail. I am confused, because the 'tree' command lists all of them under the select directory, and so far nobody can help me to delete the files and remove the directory.

G. Zarucki

If the files in question have been marked as 'hidden' or 'system', then the DOS shell won't show them, and so you can't delete them. The program's install program may have done this, to stop the files accidentally being deleted. To fix the problem, exit the shell, change to the directory in question, and type the following commands:

```
attrib -h *.*
attrib -s *.*
```

These two instructions will reset the hidden and system attributes on any files which have them set. You should then be able to delete the files in the usual way.

Local BIOS upgrade

Further to last month's piece about hard drive support in BIOS ROMs, I have just found out about a local company which can modify BIOSes to support different hard drives. The company is Custom Built Computers in Carlton, (02) 587 7104. Many thanks to Bob Barnes at RCS Radio for letting me know about them.

Second floppy controllers

I have encountered a curious problem with DOS 5.0 after upgrading from version 3.3. My system is a '386SX with four

MAY 1992

floppy drives and a hard disk. The floppy controller assigns the floppies to A: B: C: and D:, with the hard disk as E: and F:.

Under DOS 5.0, some drive re-assignment automatically takes place. The first two floppies remain as before, with the hard disk as drives C: and D:, and the last two floppies become E: and F:. This is quite handy, since it avoids the need to reassign drives for applications which assume the hard disk is C:.

The problem occurs when copying disks with PC Tools 4.3 on the E: and F: drives — copying or comparing files takes about 10 times as long under DOS 5.0 as it does under DOS 3.3. The performance of drives A: and B: is the same under both DOS versions.

However, if I run DOS 5.0's Diskcopy or Diskcomp on the troublesome

drives, the processes occur at what appears to be the normal speed. Even more curious is the fact that if I then start up PC Tools and perform a disk copy or disk compare on the drives on which DOS 5.0's Diskcopy or Diskcomp has been run, or partially run, it runs at the normal speed.

I have tried loading PC Tools using DOS 5.0's loadfix command, but it does not help. Is this one of the rumoured bugs in DOS 5.0? Can you suggest a method to overcome the problem?

Robert Allan

I use a home-grown accounting program which — deliberately — keeps all data for each of several organisations of which I am honorary treasurer, on dedicated 5.25-inch floppies, which I back up after each

session. For this reason, I equipped my computer with an extra 1.2Mb drive — the same as the A: drive.

Recently, I decided to install a third drive — a 1.44Mb 3.5-inch unit, and this necessitated the purchase of a separate 4-floppy controller board. My hard drive is a 40Mb Quantum IDE, and the floppy portion of its controller can be disabled, fortunately.

The problem is that access to the floppies has slowed down to a crawl. The BIOS ROM on the new board branded 'Unique' — appears to have to read the whole disk directory every time the disk is accessed, and during this time not even the mouse is allowed to function. Is this a normal situation, or should I have bought a different controller? I see Microgram has a combination board

Tech Tip of the Month — Counting in batch files

RECENTLY, I had the need to label groups of 16 diskettes with a volume label consisting of a constant name together with a sequential number from 1 to 16. The following batch files are my solution, and may be of interest to others as a method of sequentially counting in batch files.

```
rem VLABEL.BAT
@echo off
if %1x==x goto error
for %%a in (1 2 3 4 5 6 7 8 9 10 11 12
           13 14 15 16) do call %1 %%a
goto end
:error
echo Usage:%0 mct or % mcs
:end
```

```
rem MCT.BAT
@echo off
echo Insert disk number %1 and
pause
label b:mct%1
```

```
rem MCS.BAT
@echo off
echo Insert disk number %1 and
pause
label b:mcs%1
```

The reason for calling additional batch files from the main file was to allow for the pause, so that disks could be swapped, as label.com does not wait when the label is given on the command line, and the for-in-do loop can only have one command. I trust that this is of some interest.

I also have a question — when pre-formatted disks are purchased (Verbatim 360Kb), why does DOS 3.3's diskcopy accept them, but DOS 5.0's diskcopy insist on reformatting them during the copy process?

Ross A Young

Thanks for that tip, Ross. It is a good illustration of how to use the for-in-do loop in DOS, a construct which can be used for a lot of useful things in batch files. By the way, for those who haven't seen it before, %0 (the 'zeroth' argument to the command), is the name of the command you typed in — in this case, the name of the batch file. It's more often found in operating systems like Unix, which allow one physical file to be linked to different file names, possibly in different directories. It allows the batch file to determine which command was actually run. It also means that the 'usage' error message gives the correct name of the batch file, if somebody renames it to something else.

To answer your question — I don't really know. It seems that the formatting information is somehow non-standard, and that DOS 5.0 is fussier than 3.3 when it comes to checking the format of the disk. This problem is no doubt related to the question elsewhere in 'Tech Tips' this month about volume labels, which pre-formatted disks appear not to have.

for SCSI, IDE, and four floppy drives — perhaps this would be the answer.

I would also like your opinion on whether a suspect maths co-processor would be likely to damage the motherboard if, as I fear, it has an internal short circuit. It was installed in a previous motherboard which suffered a reversal of voltage and had to be replaced. I have not since been game enough to try it.

Robert H Abel

I've answered these two letters together because the main question in each of them relates to virtually the same problem, and is a problem which crops up frequently in relation to aftermarket floppy disk controllers. It seems that many such controllers leave a lot to be desired with regard to the on-board BIOS extension.

John Hepworth, one of our regular contributors, tried out a number of these boards when he wanted to add an extra floppy drive to his machine, and found only one which performed satisfactorily — Sysgen, available from Guardian Data, (02) 406 6144. This particular board is quite flexible in its implementation, allowing up to four extra floppy drives to be connected to the machine.

I can't say why the drives run slowly, or run slowly some of the time, but not at other times, but it's bound to be a problem in the BIOS. Whether or not the manufacturer will admit to this is another thing, but it's not likely to be a problem in DOS. The best solution when buying any new hardware is to make sure you get some sort of satisfaction guarantee, so if it doesn't work properly, you can get a full refund.

The reason that the mouse doesn't work is that interrupts are disabled during DMA disk accesses, since the disk controller and not the main processor has control of the system memory. Mind you, the amount of time that interrupts are

disabled should not be visible to the user — another sign that there is something wrong with the controller's firmware.

Finally, I'd be wary of using the fried maths co-processor in another (good), computer, although there is a good chance that it wouldn't do any damage, I wouldn't want to risk it. If it is really fried, and it didn't shut down the power supply because of an internal short, it is likely to short some of the pins of the main processor to ground or to the 5-volt rail, which may damage the processor. Also, I'd say that it's very unlikely that the co-processor survived having the power applied backwards anyway.

Missing file

I have a '386 compatible, and one of the first things I bought for the computer was Windows version 3.0. After about a week I had it up and running, except for one problem. When I tried to run it in Enhanced mode, the program would come up with a message telling me that I should have a file called 'WINA20.386' in my root directory.

The setup program didn't put this file on my hard drive, so I checked the original file, and it was not to be found. Then I found a file called WINAO386.MOD, so I decided to rename it to WINA20.386, and see what would happen. Well, it worked, and now I can use the menu, but I still think something is wrong. Programs won't run from File Manager or Program Manager. The symptom is the system just freezing — no 'Unrecoverable Application Error' messages. I hope you can help.

Richard Orr

The missing file is actually shipped with DOS 5.0, not Windows, and is required if you want to use Windows in Enhanced mode with this version of DOS.

I don't know what that other file is that you renamed, but it is undoubtedly the cause of your system crashing. You should find the WINA20.386 file on one of your DOS 5 distribution disks. Just copy it to the root directory of your hard disk.

Colour Word Processor

In response to Andrew Watts' query in the March issue relating to colour-capable word processors, your readers (including Andrew), might be interested to know that a word processor that I use at home might be suitable.

The package is called OkEditor, and I got my last update a couple of years ago from Pricom, (02) 232 4831, fax (02) 232 4355.

I purchased the word processor primarily for its ability to include technical symbols, including graphic renditions of organic chemical formulae.

I purchased it prior to acquiring a colour printer — an NEC P5XL — and have found that the colour support is excellent.

The word processor is WYSIWYG with justification, colour, underlining, italics, and graphics symbols appearing on the screen as you type.

I've found with all colour printer applications (including OkEditor), that there is not much consistency between on-screen colour and printed colour. I've just had to memorise what each package will produce on the combinations on the keyboard.

OkEditor is quite a temperamental word processor, however the latest upgrade (which I don't have), is claimed to have 'standardised' the usage of key combinations on the keyboard.

I hope this is useful for either Andrew, or any of my co-readers, who may also take some comfort in knowing that the software is an Australian product.

Bill Anderson

RGB monitor for EGA

I purchased a copy of your January, 1991 issue from a local newsstand, and was particularly interested in the answer given in your 'Tech Tips' column to a query about using VGA on an old monitor. I have a Sony KX-14CP1 RGB monitor, which I use now only for video, because I prefer a monochrome screen with a Hercules card to the CGA card I have.

Some time ago I bought an ATI EGA Wonder card from the States, but could not get it to work with the monitor. I concluded from this that it would not run EGA or VGA. Your article has given me new hope.

The KX-14VP1 has an 'analog multi' input socket with 15 pins. Is it possible to configure the pin connections as you show them for the 9-pin MultiSync

input? In Fiji technical information is hard to come by, so your help would be appreciated.

EG Newing

I am not familiar with the particular monitor which you have, but I'd say it is unlikely that it will work with a VGA card.

The parameters for composite video signals (as used by TVs and video recorders), are quite close to those for CGA colour, so it is simple for a manufacturer to build CGA compatibility into a video monitor, just by adding digital RGB inputs.

The Amiga also uses standard television scan rates, but it uses analog RGB, and it is also possible to build this into a conventional video monitor at little cost.

However, EGA and VGA signals run at much higher frequencies than TV images, due to their much higher resolution, and this requires significantly different circuitry in the monitor, as well as a higher resolution picture tube. So any monitor which is capable of displaying these video standards will no doubt make that fact well known — it adds significantly to the price of the unit.

So if your monitor doesn't say it is EGA-or VGA-compatible, then it probably isn't.

The 1991 piece in 'Tech Tips' which you mentioned applies to multi-scan or MultiSync monitors which were released before IBM ever released the VGA, but which could be used with the VGA, due to their multi-scanning nature, and the presence of analog inputs.

JUNE 1992

June 1992

Memory management

I am writing to ask if you can help in explaining the difference between the memory management techniques used in DESQview and OS/2 2.0.

I've used DESQview for a number of years, and find it very stable in multi-tasking and quite fast, although the underlying operating system is DOS. DESKview's text-based interface is functional, but it has limitations in its use of the mouse.

I would like to upgrade to OS/2, which provides the same rock-solid stability and a GUI front-end. With OS/2 there is a question about the speed of application programs.

Also, will it keep accessing the hard disk in operation, even with 4Mb of memory? As OS/2 has arrived on our shores, could you possibly present an article that compares the functionality, speed, and efficiency of both platforms? I appreciate your help in this area.

Daniel Hie

DESKview's inherent limitation lies in the fact that it is still built on top of DOS. DOS is not a very robust operating system — in fact, some people consider it nothing more than a boot-strap loader with a command interpreter tacked on. DOS was written without a thought for multi-tasking, or as it can sometimes appear, to make multi-tasking as difficult as possible.

That said, Quarterdeck has done a more than creditable job of building a multi-tasking operating system on top of DOS within the limitations that that imposes. It certainly is a lot more stable than Windows 3.0, although 3.1 is likely to give it a run for its money.

If you're finding DESQview is stable in your situation, and you don't have any plans to acquire any new applications, then my advice is to stick with DESQview. If it ain't broke, don't fix it! However, if you want to run Windows applications concurrently with DOS applications, then I'd recommend either Windows 3.1 or OS/2.

Windows, like DESQview, is constrained by the limitations of DOS, so you will always be able to run a program which can crash the entire system. OS/2, on the other hand, is a multi-tasking, multi-threading operating system from the ground up, and is pretty well protected against rogue applications. In fact, in my tests with one of the later beta versions of the operating system, I was unable to completely crash the system by running any DOS application.

This protection arises basically this is because the operating system is separated into two layers — the executive layer, which has complete access to the computer at the hardware level, and the application layer, which is where user application programs reside. These application programs can only access the hardware through calls to the operating system. Programs are free to do whatever they like to their own area of memory, but if they try to scribble over another program's memory, or the operating

system's own memory areas, the operating system will stop.

All multi-tasking operating systems impose some level of performance penalty — it takes a finite time to switch between tasks, and any processor doing more than one thing at a time will take longer than the same processor running only one task. When you're only running a single task in the foreground, with nothing in the background, the performance degradation is pretty minimal, even more so if the foreground application is running full-screen, rather than in a window.

Some brief DOS application tests that I ran on my OS/2 beta gave performance results of between 95 and 98 per cent of the figures obtained under DOS, and beta code is notoriously slow. I haven't yet received my copy of the final release version of OS/2, but it's likely to be better still.

The reason OS/2 continually accesses the hard disk is because it implements virtual memory, as does Windows when running in Enhanced Mode. This allows parts of programs to be swapped out of memory when they're not actually doing anything, so that parts of other programs can have a chance to run. Virtual memory is just a substitute for real memory, and is used to allow more programs to run simultaneously than will fit in the available physical memory.

The more physical memory your machine has, the less need there is for virtual memory. While 4Mb of RAM is adequate for most peoples' Windows needs, OS/2 is much bigger, and more memory is needed in order to attain best performance. I wouldn't recommend running OS/2 on any machine with less than 6Mb of RAM, with 8Mb a starting point for 'power users'. Also, best performance of virtual memory requires a fast hard disk, since it effectively impacts on the memory speed of the machine.

We certainly intend to keep our readers abreast of new developments in the operating system area (see the OS/2 write-up elsewhere in this issue), and we'll be giving a detailed report on OS/2 when we've had more of a chance to put it through its paces.

Stuck hard disk

A friend of mine has an IBM-compatible computer with a hard disk drive, and a 5.25 inch floppy drive. The first problem is that the hard drive will not start up when you turn the power on. The only way we can start the hard disk is by hitting the side of the computer until it starts up.

The second problem is that the floppy disk drive will not work at all. What can we do about these problems?

Shane Haddow

Your friend's hard disk is suffering from an affliction known as 'stiction', where the drive can't start up because the heads are stuck to the surface of the disks. This arises as a result of the extremely flat surfaces of the disk heads and the surface of the disk. When the disk is not spinning, the heads actually touch the surface of the disk, and the flatness of the surfaces causes a partial vacuum to be created when the disk tries to move.

Hitting the side of the drive helps the heads to break free from the disk surface, so the latter can start spinning. Once the disk is spinning, the heads float just above the disk's surface, so there is no problem. Alas, there really isn't any solution to this problem, short of getting the drive overhauled. Probably the only thing you can do is to keep hitting the drive (as gently as you can), for as long as it lives. I know a couple of people who have been doing this for over a

year, and the drives are still going strong. Just make sure you are ready for it when it does fail.

Of course, it's a bit hard to keep backups if the floppy drive doesn't work. Unfortunately, I can't really offer much help here, as the problem could lie anywhere in the drive or the controller. If the motor spins when you try to access the drive, then it could just be dirty heads — try getting a non-abrasive head cleaning diskette and running that through the drive.

If the motor doesn't start up at all, it could be a loose or broken cable. On the back of the drive there is a 34-way ribbon cable and a four-way power cable — make sure that these are securely connected to the drive. If there is a spare power cable coming from the power supply, try plugging that into the drive instead of the existing one, to see if that's the problem. Also check the ribbon cable for obvious signs that it's been pinched between something inside the computer — this can break one or more of the wires in it.

Ensure too that you haven't added any other hardware to the PC recently, which could be causing a hardware conflict. If in doubt, remove any peripherals not vital to the operation of the computer, and try again. Beyond that, there's not much you can do short of trying another drive and controller, to see where the problem is.

Logging serial data

I would like advice on what software and other accessories you could suggest to hook up an IBM XT-compatible computer to an electronic weighbridge indicator via an RS232c interface.

Ideally, data would be transmitted invisibly to the user, and would be readily incorporated into an invoicing program.

So my question is this — is it necessary to buy a set of communications functions in a library, or can programs like Supercom or Mirror, with their programmable functions, be used to automatically save data to disk as it comes from the weighbridge?

We need to capture the tare, gross, and nett weights from the weighbridge, and store them along with the data about the customer and the material supplied, which would be entered as each truck comes onto the bridge. An invoice would be printed on the spot and then posted to the customer account, so that a statement can be produced at the end of the month.

Your thoughts on the best method for achieving this task and the prices involved would be greatly appreciated. Enclosed are details of the data protocol used by the weighbridge.

Jason MacDonald

While communications programs such as the ones you mention could be used to record the data from the weighbridge, it wouldn't allow you to enter user data, such as the client's name or details of the materials weighed. You'd also have to manually correlate the tare and gross weights of each truck, so you could calculate the net weight of the material.

What you really need is a custom data-base application, which could record the weights from the weighbridge, and automatically store the gross and tare weights in the same record, along with the other information you want, each in a separate field. Setting up any data-base to store and retrieve this sort of information is a relatively trivial task using such programs as dBase IV.

However, in your case, the input is coming from two sources — the keyboard, for the client and material information, and the serial port, connected to the weighbridge. Unfortunately, dBase doesn't support inputting and outputting to I/O ports, but one third-party compiler (Quicksilver), has extra instructions to provide this access.

Using Quicksilver is probably easier than the alternative of using a more general-purpose programming language such as C or Pascal, since they are designed specifically for database operation. The only limitation is that you will be limited to 'polled' operation of the port — your program has to continually check the port for the arrival of new data, at least after you have polled the device to get it to send the data. With polled access, there is the danger of missing data if your program doesn't read the port fast enough to keep up with the incoming data stream. For this reason, I'd slow down the speed of the port as far as possible, (in your case 600 baud). The amount of data transferred each time is relatively small, so you won't notice the speed, but you'll make the timing of the program less critical.

Then, all you need is a bit of code in your program to interpret the incoming data stream, so you can filter out the actual weight data and place it in the database. All this would need to do would be to set the serial port up with the correct speed and other parameters, then send the 'read' command to the port whenever a truck is weighed, and record the weight data as it comes in.

If you were to use a general purpose programming language, you would be able to make the data handling more streamlined and faster, but you would then be faced with programming all the database functions that dBase can provide. Since the actual communication with the weighbridge is only a small part of the program's function, I'd be inclined to try the database approach first.

By using the dBase file format, you will then be able to import the collected data into any program which can read dBase files.

Quicksilver is distributed by Micro Data Management Systems, on (02) 745 1677.

Why copy protection works

Recently, I read in July 1991's 'Tech Tips', a piece on copy protection. Except in the case of password entry to a program (which can be backed up anyway), I can't see why a copy program, designed to read and copy byte for byte, can't create a disk that is identical in every way to the original.

I've been told that some companies use a laser hole in the disk, but since the program can check for the laser hole, it should be possible to have the area with the laser hole on the backup disk marked so that the read head gives the program the same messages as the laser hole.

Mark Whitaker

There are a number of tricks that software publishers use to make copying their disks more difficult, some of which are more easily circumvented than others. The laser hole method is one of the most secure, but is also pretty rare, because it requires specialised equipment to duplicate the disks. The way the software tests for the presence of the hole is to write some data to the disk where the hole should be, and then read it back. If the data reads back without error, then there's no hole there. Because of this, there is nothing that a disk copier can write to the disk to make it appear that the hole is there.

An alternative method takes advantage of the cyclic redundancy check (CRC), information recorded with each and every sector of data on the disk. The CRC is an error detection scheme which is used to check the validity of data read from the disk. However, this particular copy protection scheme records an invalid CRC with one or more sectors on the disk, so that any attempt to copy those sectors will produce a disk error.

If a copy program nevertheless copies the disk verbatim, the CRCs

JUNE 1992

I HAVE BEEN PROGRAMMING for some time now, and it was not until recently that I included mouse support in any of my programs. I believe that mice have an aura of mysticism when it comes to programming. Many programmers have written-off mouse support for their programs because not enough information is available on the subject — in fact, I had to research for some time in order to gain the information used in my programs.

This may be because I looked in all the wrong places, but when I did find the information (in a crusty

old manual called 'MS-DOS Language System Programming', written in 1986 by Olivetti), it took a book on assembly language and many exhausting nights of trial and error. I believe that you can buy a Microsoft Mouse Reference Manual, but how many people know about this? Finally, after a lot of blood, sweat and tears, I have a library of mouse support for my programs.

I have included listings in C and Basic for primitive mouse support.

```
#include <dos.h>
#include <stdio.h>

int mouse_cont (int ax1, int bx1, int cx1, int dx1);

t main ()
{
    int ax1 = 1; /* set the variables for mouse_cont */
    int bx1 = 0;
    int cx1 = 0;
    int dx1 = 0;
    mouse_cont (ax1, bx1, cx1, dx1); /* perform function call */
    return;
}

int mouse_cont (int ax1, int bx1, int cx1, int dx1)
{
    _AX = ax1; /* place values into the registers */
    _BX = bx1;
    _CX = cx1;
    _DX = dx1;
    geninterrupt (0x33); /* call the interrupt */
}
```

For the C listing I defined a function called `mouse_cont` (not to be confused with the interrupt function — explained later), and specified that it will receive four variables corresponding to the four registers. In this example, I placed the value 1 into `ax1` — this is the interrupt function that I wish to use.

For the basic listing, I use the variables %1 through %4 (programmers can name them themselves), and the example call given (from line 100 on), places the value 3 in %1 — telling the mouse driver that interrupt function 3 is required, when it is called with the 'call mouse' statement in the program. The variables %2, %3 and %4 carry the button status, and X and Y co-ordinates, respectively.

will be recalculated by the controller during the copy process, so a different CRC will be recorded for those sectors. The program knows to look for a CRC error, and will get upset if it doesn't find it. As I understand it, the CRC is actually calculated by the controller itself during write operations, so the actual value recorded is not under software control.

Fortunately, copy-protection schemes have largely fallen out of favour of late, and the ones which remain are certainly not as insidious as past efforts. One popular system simply records the name of the owner when it's installed on the hard disk, so any copies made from the hard disk will point the finger squarely at the original owner.

Some schemes also try to record

the name on the distribution disk, which I don't like — all distribution disks should be write-protected before ever going near a disk drive, to guard against virus infection. If your machine does suffer a virus attack, you should be able to go back to your original disks and re-install the software knowing that it hasn't been contaminated.

If you have software that does try

mouse support for your programs

```

10 DEFSEG = 0
20 MSEG = 256 * PEEK (51*4+3) + PEEK (51*4+2)
30 MOUSE = 256 * PEEK (51*4+1) + PEEK (51*4) + 2
40 DEF SEG = MSEG
50 REM lines 10-40 must be executed before any calls to mouse driver
100 REM get the X and Y mouse position
200 %1 = 3
300 CALL MOUSE (%1, %2, %3, %4)
400 X = %3 500 Y = %4 600 BUTTON = %2

```

The listings I have included are by no means in-depth, and require a bit of thought and at least a meagre knowledge of assembly language. However, it is not hard to include these examples in any programs you may write, and try them out by switching the interrupt functions around and watching their effects. I suggest you experiment a little, and if you want to go further, buy a suitable book on the subject. Note that these listings only work with Microsoft and compatible mouse drivers.

The Microsoft mouse driver has about 35 interrupt functions — below is a list of the more important ones. The interrupt function is the value placed in variable ax1 for the C listing, and %1 in Basic.

Function 0 — Mouse initialisation.

This returns 0 if correct mouse hardware is detected, otherwise it returns -1. AX returns the mouse status, and BX the number of buttons.

Function 1 — Show mouse cursor.

Shows the mouse cursor in whatever graphics or text mode you are currently in. Returns nothing.

Function 2 — Hide mouse cursor.

Hides the mouse cursor in the current graphics or text mode. Although hidden, you can still get its position. Returns nothing.

Function 3 — Get position and button status.

This function returns the X and Y coordinates of the mouse cursor, and the button status. CX returns the X-co-ordinate, DX the Y-co-ordinate, and BX the button status, (0 if button is up, 1 if button is down).

Tony Osborne

to write to the distribution disk during installation, make a copy (after write-protecting the original), and install from that copy.

Hardware locks, or dongles, are a popular way to protect software from unauthorised use, while not inhibiting its ability to be copied in any way. This allows the software to be installed on several machines, and backed up as desired, but still

ensures that only one machine can run the software at any one time. Dongles usually connect between the parallel port and the printer, and are supposed to be transparent to normal printer use. However, many people have found otherwise, and have installed a separate port for the dongle.

The need for copy protection has also been lessened by giving the

authorities in most countries some teeth to allow appropriate penalties to be levied against software pirates. This, and a few well-publicised convictions and out-of-court settlements, have made most companies very wary of using pirated software, and most larger corporations now carry out software audits as a matter of course, to make sure that they're operating legally.

July 1992

Boot.sys alternative

I own a recently purchased '386DX, and I single task 80 per cent of the time, so I use PowerMenu to access my applications. I run my applications concurrently 15 per cent of the time, using DESQView, and because I don't like Windows, I use it five per cent of the time. Like others in this situation I use different config.sys and autoexec.bat files, by renaming them for my different boot-up requirements.

After reading about Boot.sys in the November 1991 issue of Your Computer, I purchased a copy to trial because here seemed to be the answer to my boot-up problems, only to find that Boot.sys is not compatible with my operating system — DR DOS.

Do you know of any alternative software that is compatible with DR DOS which will do the same job? I have written to Budgetware suggesting that they give clear warning of the non-compatibility of Boot.sys with DR DOS in their catalogue.

Colin McIntyre

For starters, if Budgetware didn't advise you of the incompatibility when you bought the software, they should at least refund your money. That doesn't help with your problem, but at least you haven't wasted your money on a product that you can't use.

Fortunately, DR DOS 6.0 has a couple of features built-in, which you can use to achieve virtually the same result as with Boot.sys, for no extra outlay (aside from a few minutes at the keyboard). This is due to the addition of several useful commands in the config.sys file which MS-DOS does not have, which really go to make Boot.sys unnecessary with DR DOS.

DR DOS has a command called 'switch', which is a bit like the 'case' command in Pascal, allowing execution to conditionally branch to one of a number of optional

subroutines. You can use the 'echo' command to print a list of available options on the screen before using the switch command. With DR DOS, you can also set environment variables from within the config.sys file, allowing information about which option was selected to be passed on to the autoexec.bat file.

Here is a sample config.sys file which should get you started:

```
timeout = 5
echo = Press 1 for Power Menu
echo = Press 2 for DESQview
echo = Press 3 for Windows
switch menu, dv, win
exit
:menu
rem DESQview selected
device=qemm.sys
set OPTION=2
return
:win
rem Windows selected
set OPTION=3
device=himem.sys
return
```

The 'switch' statement calls sub-routines (hence the 'return' statements), and the 'timeout' statement causes the first option to be selected automatically if you don't select an option within five seconds. Then, all you have to do in your autoexec.bat file is to test the value of the 'option' environment variable, to see which key was pressed.

Music Programming

I have a question and a tip to do with music. I intend to do some programming for music teachers, and am not sure which programming language to use. The software will be used with MIDI keyboards, and involve music notation, sound through the keyboard or internal chip, and graphics and text display. I also want the

programs to be easily portable between IBM, Mac, and Atari platforms.

My tip is for those people who wish to play music through a MIDI system. It is not necessary to purchase a MIDI keyboard — an alternative is to buy a sound module or tone generator, (as long as it is multi-timbral). Examples are the Roland U220 and the new half-rack Sound Canvas. Yamaha is coming out with one soon for \$695, without the need for an interface. For the budget-minded, the old Roland MT32 modules, second-hand, are around \$250, still with much software backup for this great little device.

Rupert Sherwood

Your music tip is probably more applicable to Eric Holroyd's 'Take Note' column, but it is appreciated nonetheless. The only problem with not having a keyboard is that it precludes real-time programming, but if all your programming is going to be step-time, then you won't notice it. Of course, it also serves as a backup for the PC, for when you want to do something the 'old-fashioned way'.

As for writing music software, the term 'easily portable' is something of an oxymoron. The reason being that in the three machines you mention, you are dealing with completely different hardware, and completely different software. While your core code might be the same in all three cases, any operations that require keyboard input, screen output, disk access or access to the MIDI port, are likely to have to be programmed differently for each machine.

As you can imagine, this is bound to be a significant part of your programming effort. As an example, the people at Aldus tell me that 60 per cent of the programming in PageMaker is common to both the Mac and Windows platforms, with the other 40 per cent being machine-specific. So while it is possible to create portable code, it is rarely easy.

JULY 1992

Even programs written for operating systems like Unix, which spans a huge number of different platforms, often have to be fiddled with when they are ported from one machine to another.

As for choice of programming language, I'd opt for C, or even better, C++. The object oriented capabilities of C++ would I think be useful for your application, especially for graphics handling. C was originally developed as a systems programming language (for writing the Unix operating system, incidentally), but has rapidly gained acceptance as an application programming language too, and, with the C++ extensions, is finding even more favour among application programmers.

C code is also pretty portable, due to the wide availability of consistent C implementations for a wide range of platforms. Whatever your choice for a programming language, it is imperative that you can obtain a compiler for each platform which behaves in much the same way as the others. The whole point of portability is to minimise the amount of code which has to be re-written for each machine, and having consistent compiler command sets is the first step.

Upgrade or renew?

I currently own a PC 10-50 XT-compatible, and I bought it because I knew it would be useful in primary school. As I entered High School, and started to see what we do at school with IBM computers (word processing, spreadsheets, DTP, and CAD), I realised that my computer didn't have enough memory for those jobs.

I consulted a computer expert, and he advised me that it wasn't worth upgrading or adding more memory or better

hardware. He said I should seek out information about what computer would suit my needs. That's why I buy your magazine and look for any hints or tips on buying a computer.

Could you tell me what computer would suit me for the above tasks?

Luis Tafigliola

I have to agree with your 'computer expert' — upgrading an XT to something that can be used with today's software is not usually worth the hassle. By the time you've bought a new motherboard, memory, a hard disk and controller, and a high density floppy drive or two, you will probably have spent about as much money as you would to buy a new computer, and have a lot of left-over bits from your old computer that you can't use.

I think a much better approach is to forget your current machine (either keep it or sell it), and start with a new machine. At a minimum, I'd recommend a '386SX-based computer with at least 2Mb of RAM, and a 40- or 60Mb hard drive. That will run most any application, including Windows and applications which only run under that environment. While there are still '286s around, they are a dying breed, and more and more software requires the capabilities of a '386 or better chip.

For basic desktop publishing or CAD, this machine would be adequate, but if your plans in either of these areas are at all ambitious, I'd look seriously at twice that much memory, and even at a '386DX or '486SX machine, if you can find one at the right price. While a '386DX may not be much faster than an SX with DOS and Windows programs, the new generation of operating systems will run a lot faster on a DX.

Usually, upgrading the processor means replacing the motherboard, which is a major component of the system, and is

often not economical. There are a few machines around which do allow for plug-in processor upgrades, but these tend to be in a pretty high price bracket, so even their economy is doubtful.

(For more information about buying a PC, see Jake's article, 'A Painless Approach to Buying a PC', in the May '92 issue.)

Two IDE drives

I have a 42Mb IDE hard drive, and I wish to upgrade to a larger one. From reading your magazine and other publications, it seems that one can't mix and match these drives. If you need a second hard drive then it must be of the same type and manufacture, and you must use it in slave mode, otherwise strange things start happening.

My hard drive uses interrupt 14, so my question is this:

Suppose I have a spare interrupt and a spare slot, and modify the card that plugs into the PC's slot so that it uses a different interrupt. Would I be able to use my old drive on, say, interrupt 13, while the new one uses interrupt 14?

Each drive would have its own card and interrupt. The card on my existing drive is capable of having the floppy controller on it disabled. I looked at the card, and if I cut the track on interrupt 14, then I can patch it to any interrupt that is available.

So to my question — will it work? I don't see why not, but then maybe I'm missing something.

WJ Hughes

Unfortunately, adding a second hard drive is not as simple as just changing the interrupt. While there is no physical reason why you can't have a hard drive using another interrupt, DOS and the BIOS will be expecting it to use interrupt 14, so at best, it won't work.

JULY 1992

ONCE UPON A TIME, Microsoft made a run-time version of Windows available, which allowed a single Windows program to be run, without having the full-blown copy of Windows on your hard disk. With Windows 3.0, this ceased to be the case — now you have to have Windows on your machine in order to run a Windows application. However, this tip is handy for people who want to be able to run a single Windows application, while keeping the full Windows environment hidden.

me

Small children use my PC at home, and I like to keep them away from sensitive areas such as the Windows Desktop and File Manager, because I don't want my configuration modified. At the same time, I would like them to have access to some programs which only run under Windows. So far, I have found no way to password-protect Windows applications on the Desktop, though passwording is possible in the DOS 5 shell.

The procedure described below allows a user to access a Windows application from the DOS 5 shell,

without any opportunity to enter the main Windows environment, by temporarily substituting the required application for the Program Manager as the Windows shell.

The system.ini file contains the following lines, which define the shell which controls Windows:

```
[boot]
shell=progman.exe
.....(etc)
```

With this configuration, Windows stays active until Program Manager is closed. But any other program can be substituted in the 'shell=' line, so that Windows not only starts in that application, but also shuts itself down when it terminates. For example, changing the line as follows will cause Windows to start Windows Write, and return to the calling environment when Write is exited:

```
shell=write.exe
```

Note that if the executable is not in the path, the full pathname must be specified.

Because I also want to use Windows normally, I don't want system.ini altered permanently, so I wrote a batch file to handle the modifications and restoration. As the batch file is used for multiple applications, it requires a single parameter to be passed on the command line; and for each possible parameter, a copy of system.ini must be created as a template, and modified as above, and named to suit the batch file — for example, system.bg. A template of the 'normal' system.ini must also be created by copying it to sysini.std.

(System.ini is normally only 1300-1400 bytes long, so the creation of multiple templates is not a major problem in terms of disk space, as each will live within one allocation unit.)

To use the batch file, a program item is created in the DOS Shell for each application, with a command line like the following models:

```
c:\bats\qikexwin.bat bg
or
c:\bats\qikexwin.bat wr
```

There's another consideration — that of the I/O address of the controller (which is actually on the drive itself, not on the expansion card). If you were to put two drives in the same machine, which, with IDE drives, means two controllers as well, they will both try to respond to the same addresses, and neither will work. While there is a designated secondary address for another hard disk controller, again, the software has to be modified to work with this address.

The master/slave relationship of IDE drives is designed to get around this, allowing the two drives to share the same set of I/O addresses and the interrupt line. Unfortunately, as is often the case

with new standards, a few early implementations of the IDE standard were mutually incompatible, this limiting master/slave systems to ones in which both drives were from the same manufacturer.

Ideally, you can take any two IDE drives, and set one to operate as the master, and the other as slave, and have them appear to the BIOS and DOS as your first and second hard disks. If your drive is one of those early problematic ones, then your alternatives are either to try to locate another drive of the same make and model, and use that as the slave drive, or get a single new drive and abandon your old one.

But before deciding that it's not

going to work, obtain a new drive (from somebody who will refund your money if it doesn't work), and try it out. Be sure to try both combinations of master and slave if one doesn't work — there is a good chance that it could work one way but not the other.

Don't overlook the software alternatives either — unless you have a lot of already-compressed data on your drive, one of the popular software disk compression packages reviewed elsewhere in this issue might be suitable for your needs. While they do slow down the system a bit, and use up some memory, if you can live with this they are cheaper than a new hard disk.

Run time Windows

To further protect the system from inexperienced users, the 'Shell to DOS' option in the DOS Shell is also passworded. The DOS Shell itself is run from a special batch file called by autoexec.bat, which pre-

vents the command line from being accessed through Alt-F4. But that is another story.

Finally, this technique would be useful if you normally don't work in Windows, but use one of its applica-

tions fairly frequently — for example, the calculator, and want to avoid the need to go through several mouse-clicks and dialog boxes to end it.

John Carroll

@echo off

REM QIKEXWIN.BAT - to run a Windows application

REM without using Program Manager - J Carroll 1/5/92

REM Go to message if no parameter

REM was passed to batchfile

if %1x == x goto :BADPARAM

goto %1

REM parameter was "bg" to run Backgammon.

:bg

echo Loading Backgammon...

goto :FORALL

REM parameter was "wr" to run Write

:wr

echo loading Windows Write...

goto :FORALL

REM Put labels and command blocks for any

REM other parameters here

:BADPARAM

REM batch file will fall through to here

REM if parameter incorrect

echo Incorrect syntax or bad parameter - check before
retrying.

echo Allowable parameters are: bg wr

pause

goto :ENDIT

:FORALL

REM replace system.ini with version modified to use the

REM executable instead of progman.exe as the shell

copy c:\windows\sysini.%1c:\windows\system.ini > nul

REM kill the parameter if it exists

shift

REM Windows will run the required

REM executable automatically.

win

echo Unloading Windows...

REM Restore system.ini to normal

copy c:\windows\sysini.std c:\windows\system.ini > nul

:ENDIT

Locking the hard disk

I am looking for a hardware method to isolate the hard disk in my computer, so I can use the computer without catching viruses. It is a 20MHz '386SX with 3Mb of RAM, a 42Mb Seagate ST151 hard drive, 1.2Mb and 1.44Mb floppy drives.

I could fit a switch to the four-wire power supply cable to the hard disk, but I'm worried that if the controller card has higher voltages on inputs and a disconnected supply rail, I will damage the controller. I want to be able to use the machine with its attached printers, modem, and LAN, with absolutely no chance of viral infection, so that other people can use my machine safely.

I was infected with the Michelangelo virus by a computer consultant who should have known better. The keyboard lock and virus protection software is no use when someone starts from an infected floppy.

Andrew Miller

If you really want to disconnect your hard drive from the rest of the system, then I suggest one of the removable hard disk carriers sold by, say, Rod Irving Electronics. These house the hard disk in a removable caddy, which slides into a receptacle mounted in the PC. I certainly agree that disconnecting the power while leaving the signal cable connected is a bad idea, for the reason you cite.

However, there is another issue, in that the disk controller won't be able to find the drive when you boot up the computer (you shouldn't insert or remove the drive with the power connected), and it will make the boot-up process rather lengthy, while the BIOS tries to access the drive. This is likely to make booting up the computer without the hard drive installed more of a hassle than it's worth.

Since you're obviously not satisfied with the level of protection offered by virus-protection software, I suggest you look at a hardware protection solution, such as the Thunderbyte board from Calmer Utilities, (02) 487 1715. This installs itself as an extension to the

AUG 1992

BIOS, and so can't be overridden by a simple act like booting from a floppy. That would leave your hard disk protected from viruses, without exposing it to potential damage by removing it.

Dirty heads

After having recently upgraded from my old XT to a '486 PC, I was going through some of my old 5.25-inch disks, when I started getting the old 'Abort, Retry, Fail, Ignore?' message. As it was getting late at night I did not try to find out what the problem was at that stage.

A few days later I was trying to install some new software which came with

only 5.25-inch disks, when the same messages started appearing. I could use the 'dir' command and see what files were on the disk, but I could not access or execute any of the files. The thought of dragging the new computer back into the city just to get the floppy disk drive looked at certainly wasn't comforting.

So I decided that if I was going to take the PC in, I may as well make sure the drive was well and truly broken. I pulled the floppy disk drive out, removed the cover, and discovered that there was a blob of magnetic oxide sticking to one of the heads. Cleaning with hydrocarbon spray and careful replacement of the drive was all that was required to bring life back to my floppy disk drive. After digging out the old disk that I was trying to read the previous day, I noted that it was one of those no-name disks, and indeed the first few tracks were missing their oxide layer

which was obviously attached to my drive heads!

So the moral of the story is — don't try to read old, dusty, no-name disks, which are not really required urgently, and if your floppy drive suddenly refuses to read 'good' disks anymore, then try a head cleaning disk or solvent before you race back to the computer dealer.

Michael Roe

There has always been a lingering thought in the back of my mind about some of those cheap disks, and the quality (or otherwise), of the binder used to stick the oxide to the surface of the disk. While they may appear to work well at first, the possibility of the oxide coming adrift over time is a very real one.

August 1992

The hazards of caching

I recently had a problem with a friend's computer, which surfaced when I tried to run a program that I knew to be on the disk in his 5.25 inch drive. It came up with a 'file not found' error. Upon getting a directory of the disk, I found it to be an exact mirror image of the last disk that I had used in the drive.

Had one disk been copied over the other? Had a new virus written the directory from the last disk onto this one? Fortunately, it was nothing of the sort. Checking the disk in my own machine, I found it to contain exactly the files that it used to, all of which loaded and executed correctly.

My friend put this down to a hardware fault — a damaged drive cable. Not satisfied with this explanation, we checked his machine and found the culprit — a badly-written disk caching program. It couldn't tell when a drive door was opened. It either could not read the status of the drive, or didn't bother to.

I have a personal hate of many disk caching programs, as they do not allow you

to tell what has been written and what is still waiting to be written to the disk. I do run a cache myself — SmartDrive from Windows 3.0. This I find to do the job well, even though Windows itself has been taken off the hard disk and put back on the shelf months ago.

Many people use drive caches without knowing the dangers. When they cause failures people do not put them down to the drive cache. The biggest problem is delayed-write caches — as well as reading in the information expected to be used next, they have a feature to hold data in memory and writing it disk when the computer is idle.

This speeds up disk accessing enormously, but it does come with a cache — if the data hasn't been written to the disk and the computer crashes, the data is lost. As well as that, there may be files left open on the drive which may or may not be recoverable. The version of SmartDrive that comes with Windows 3.1 works this way. If you wish to run the risk of the computer losing data when the computer crashes, then a read/write cache is the way to go. But for my valuable information, I'll stick to a read-only cache.

Scott Neville

It sounds like the disk cache that your friend was using was written exclusively for use with hard disks, and didn't bother to check the media had changed between successive operations — an impossibility with a conventional hard disk. Still, in that case, it should refuse to cache a removable drive, if it cannot cope with the disk being changed.

As far as delayed-write (or write-back), caches go, I am willing to live with the slightly increased risk of data loss, and benefit from the greatly improved performance that comes with such a cache. Remember, you're going to lose data if the computer crashes, (unless you're lucky enough to experience the crash immediately after having saved your work), whether you're running a delayed-write cache or not.

Crashes are an unfortunate fact of life for many people, but if they happen more than occasionally, then a pretty fundamental problem in either the software, or hardware, or both, is indicated.

Delaying the write operation only with a cache only slightly increases the window of time during which a crash will result in loss of data. You're not magically immunised against data loss by abstaining from using a cache, merely reducing the probability by a few per cent, if that.

Some years ago, I had the dubious privilege of working on a computer which crashed frequently — as often as several times an hour, on a bad day. This rapidly instilled in me the instinct to save my work frequently — a reflex which has remained with me to this very day. Frequent saving of your work is the best prevention against losing a serious amount of your work, and it doesn't help matters when industry leaders (such as Microsoft in most of its Windows applications), make saving a chore requiring three or more keystrokes. I've had 'save' on my F2 key for as long as I can remember, and whenever I pause to think, I hit it instinctively.

Caching floppy drives is a different matter — with write-back caching, there is a very real chance of removing the disk before the data has been written to it, under the impression that because you have saved, the data has been written. For this reason, Windows 3.1's SmartDrive doesn't use write-caching on removable disks, although it can be enabled if you really want to risk it. Also, no form of caching should be used on network drives, unless it's done on the server itself. Somebody could modify a file that's cached on your machine, and you could end up using old data without realising it.

If you don't like write caching, you can disable it when you load SmartDrive by specifying your hard drive names explicitly, and not putting a '+' after them. If you do use write caching, remember to flush the cache before turning the machine off, just to make sure all

the data has been written to the disk. A little batch file will make this easier, or you could add it to your 'park' program, if you have a stepper motor drive which needs parking.

Like many things, a write-back cache has a lot going for it, if you know what the dangers are, and how to minimise them.

Printer information request

I have a problem perhaps you or one of your other readers may be able to help me with. After assisting a man with a physical task, and declining payment for my help, he gave me a printer for my computer. The printer is a Canon colour printer, model A1210, but unfortunately he had lost the manual for it. The printer appears to have been well looked after, and I would love to be able to use it, and this is where my problem comes on the scene.

I phoned Canon in Sydney, seeking information on this machine, and all they could tell me was that it was introduced in 1983, and that they don't have any spare parts for it. They didn't know if it used any special type of paper (such as rolled paper), or if there was any substitute driver that could be used with it. However, they thought the liquid ink cartridge used on this machine would be interchangeable with current Canon cartridges.

So now I have an 'as new' colour printer, and no way of using the machine. Could you, or any of your readers, assist me with getting this printer up and running? Perhaps a subscriber has an A1210 printer, or may have owned one, and can assist with any information. I would be most grateful for any assistance.

George Forbes

I don't have any information on the printer model you mention, but if any readers out there have any data on this printer, we'll gladly pass it on to Mr Forbes.

Monitor upgrade

I am a great admirer of your magazine. Ever since I have started delving into computers I have been getting tired of seeing my plain 4-colour monitor in front of me. Apart from buying a new monitor altogether, I have heard talk of using VGA cards with older monitors.

I have asked around, but no one seems to know the details of what you need to use one, and what the result will be. Do you have any suggestions?

Richard Nichols

By your description, it sounds like you have a CGA monitor, which unfortunately, is unable to cope with VGA signals.

There are two reasons for this — the resolution of VGA is much higher than that of CGA, and the CGA monitor is not able to cope with the higher frequencies involved.

Also, VGA uses analog signals for the colour information, which makes the monitor capable of displaying any number of colours, limited only by the capabilities of the card. CGA, on the other hand, uses digital colour signals, so the monitor can only display the number of colours that it was designed to — 16 in total.

Some pre-VGA multi-frequency monitors were able to work with VGA cards, as they had both digital and analog inputs. These monitors were designed to work with EGA and Super EGA cards, and had no trouble coping with the slightly higher frequencies of the VGA and Super VGA signals.

I used an original NEC MultiSync with a VGA card for several years in this way, with no problems whatsoever.

Unfortunately for you, such a technique is not applicable to CGA monitors, (or standard EGA monitors, for that matter).

AUG 1992

Monitor loses colour

AUG 1992

I use a 10MHz '286 with an Acer VGA colour monitor, and 256Kb video card from DSE. My problem is that on occasion, the monitor will start up in black and white.

The quirk is that word processing and accounting packages will appear in monochrome, and stay this way for several reboots. Although, when this occurs, monochrome games such as *Civilisation* and graphics packages like *Dr Halo* will come up in colour.

This pretty much ruled out a monitor fault, so I checked the setup data, the DIP switches in the main unit, and even had the monitor cords cleaned just in case. None of this worked — do you have any ideas?

My other question relates to monitor care — is it a good idea to leave the monitor permanently switched on? If so, is it recommended to leave the rest of the system running in order to use a screen blanker, or just leave the monitor switched on with the brightness turned down?

Adam Burberry

The problem you are having with your monitor coming up in monochrome instead of colour is a relatively common one, and is a result of the lines in the monitor cable which tell the VGA card what sort of monitor is attached. While some monitors simply tie to ground the line which indicates the presence of a colour monitor, others seem to have active circuitry which does this when the monitor is turned on.

If the monitor takes too long to power-up when switched on, or is switched on after the computer, then the VGA card may not see that line as grounded, and assume the presence of a monochrome monitor. This check seems to happen on

ONE OF THE biggest problems in writing batch files is the lack of any interaction with the user beyond the command line — there's no way built into DOS to get a response from the user. These four short programs give batch files the ability to read a single key press. It should greatly simplify the writing of batch files such as menus.

`Get_key.com` is the shortest and simplest of the four — it takes up only nine bytes, and contains only four commands. `Get_key` sets the DOS errorlevel to the ASCII code of whichever key was pressed. `YesNo.com` sets the errorlevel to 1 if the user pressed the 'Y' key, to 0 if he pressed 'N', and to 255 if anything else was pressed. `Get_num.com` sets the errorlevel to the numeric value of the number key pressed, or to 255 if the key wasn't a number. `Get_alfa.com` saves the batch file writer having to worry about checking upper- and lower-case values for letters, it sets the errorlevel to 1 if the 'A' key was pressed, 2 for the 'B' key, through to 26 for the 'Z' key. If the key was not a letter, then it's set to 255.

Listing 1 is the assembly language source code for `Get_num`, and shows the basic operation of all four programs. `Get_num` begins by calling interrupt 21 (hex), function 1. This waits for a key to be pressed,

some cards only when the card is first powered up (along with the rest of the machine), so rebooting doesn't help.

This is then conveyed to the operating system, which then does not attempt to display any colour information.

Some software doesn't read the display type from the operating system, but has it set internally, and would be unaware of DOS' assumption that you only have a monochrome monitor, and display the

and puts its ASCII code into the AL register. The next section does something with that value. In this case it checks if the value is between the code for ASCII '0', and ASCII '9'. If so, the the code for '0' is subtracted, leaving a value between 0 and 9 in AL, otherwise 255 is put in AL.

The final section of the program uses interrupt 21 function 4C, which ends the program and sets the errorlevel to the value in AL.

Listing 2 is a Basic program which will make `Get_key.com` when it's run. By substituting the different program names and lines of code in Listing 3, the other programs can be written. I've used this method because not everybody has an assembler, but everyone with DOS has a form of Basic.

Otto de Leeuw

colour information as it would do normally.

The simplest solution is to ensure that the monitor is switched on before the system unit.

This can be difficult though, if the monitor draws its power from the system unit. Then the alternative would be to jumper pins 10 and 11 in the monitor's VGA plug, thus forcing the card to recognise a colour monitor.

I can't see any reason why you should leave a monitor

the Month — Interactive batch files

Listing 1:

```

CSEG SEGMENT
    ASSUME CS:CSEG,DS:CSEG
    ORG 100H
MAIN PROC NEAR
    MOV  AX,700H      ; Wait for key press,
    INT  21H          ; put ASCII code into AL
    CMP  AL,'0'        ; Is AL a numeric char?
    JB   Error
    CMP  AL,'9'
    JA   Error
    SUB  AL,'0'        ; If so, subtract ASCII 0
    JMP  Exit
Error: MOV  AL,OFFH    ; If not, put 255 in AL
Exit:  MOV  AH,4CH     ; Terminate with errorlevel in AL
    INT  21H
MAIN ENDP
CSEG ENDS
    END

```

Listing 2

```

100 REM Basic program to create Get_key.com
110 CLS:PRINT "Creating Get_key.com"
120 OPEN "GET_KEY.COM" as #1 LEN = 1
130 FIELD #1,1 AS A$: TOTALSUM# = 0
140 READ LINES,LENGTH, PROGSIZE#
150 FOR I=1 TO LINES
160 FOR J=1 TO LENGTH: READ B$: TOTALSUM#=TOTAL
    SUM#+VAL("&H"+B$)
170 LINESUM#=LINESUM#+VAL("&H"<170>+B$)
180 LSET A$=CHR$(VAL("&H"<170>+B$)):PUT #1
190 NEXT J

```

```

200 READ LINETOT$:LINECHECK#=VAL("&H"+LINETOT$)
210 IF LINECHECK#=LINESUM# THEN 230
220 PRINT "ERROR IN LINE #";270+10*I
230 NEXT I
240 CLOSE: IF TOTALSUM#<<>>PROGSIZE# THEN 260
250 PRINT "Get_key.com created successfully":SYSTEM
260 PRINT "Get_key.com is invalid":SYSTEM
270 DATA 1,9,923
280 DATA B8,00,07,CD,21,B4,4C,CD,21,39B

```

Listing 3

To make the three other programs change Get_key.com everywhere it appears in Listing 2, to the correct name below, and replace lines 270 onwards with those below.

YesNo.com

```

270 DATA 2,19,3343
280 DATA b8,00,07,cd,21,3c,59,74,12,3c,79,74,0e,3c,4e,74,
    10,3c,6e,5b7
290 DATA 74,0c,b8,ff,4c,e9,09,00,b8,01,4c,e9,03,00,b8,00,
    4c,cd,21,758

```

Get_alfa.com

```

270 DATA 2,19,3350
280 DATA b8,00,07,cd,21,3c,41,72,16,3c,5a,77,05,2c,40,
    e9,0f,00,3c,564
290 DATA 61,72,09,3c,7a,77,05,2c,60,e9,02,00,b0,ff,b4,4c,
    90,cd,21,7b2

```

Get_num.com

```

270 DATA 2,12,2153
280 DATA b8,00,07,cd,21,3c,30,72,09,3c,39,77,380
290 DATA 05,2c,30,e9,02,00,b0,ff,b4,4c,cd,21,4e9

```

switched on when not in use, except that it eliminates the power surge which results when the monitor is next turned on.

Any well-designed monitor will have components to limit the sudden rush of current when power is first applied, and by turning off the monitor you're saving wear and tear on the rest of the components in the circuit. Besides, if the power supply does eventually fail, the repair would typically be a lot cheaper than one to

the main circuitry of the unit.

Many people leave their computer running all the time (BBSes, files servers, and the like), but the monitor off when unused, to save wear and tear.

A screen saver only saves the phosphor coating on the screen itself — the rest of the monitor still runs full-tilt displaying a blank screen, (or little fishies swimming around).

Oh yes, it also saves power, which is something we should all be trying to do.

More memory for an XT

I have an XT turbo running at 8.5MHz, with the usual 640Kb of RAM, and a 20Mb hard disk. Nowadays, I am finding that my old XT is struggling with today's memory-hungry software. The computer has several TSRs installed at bootup time — a disk cache, an ANSI driver, a typematic rate accelerator, and a mouse driver. Through experimentations with DOS 4.01's install command in the config.sys file, I have managed to free up a

AUG 1992

little bit of memory, leaving 510Kb free. This is hardly a magnificent amount, considering my other computer (a '386 machine running DOS 5.0) has something like 620Kb free, with all of its TSRs loaded high. Through the use of the Manifest program, I have discovered that my XT has 248Kb of free memory above the conventional 640Kb!

Memory management programs, such as QEMM, don't seem to like working on my XT, saying that the machine doesn't have the appropriate memory expansion card to access this high memory, and that it is not an 80386-class machine.

Shane Lim

The message you are getting from QEMM is not 100 per cent correct — what it should have said is that you have 248Kb of free memory space. However, in an XT, there is no memory occupying these free addresses — it's just empty nothingness as far as the processor is concerned. It's true that this space is also empty in ATs and '386s, but it is possible to map unused extended memory into

these gaps, by using the memory management features of the NEAT chipset (present in some '286s), or built into the '386 chip itself.

This hardware support is necessary for programs like QEMM or QRAM (its '286 sibling) to be able to load programs in here — loading programs into non-existent memory is doomed from the outset. Both processors can also address the bottom 64Kb of extended memory while in real mode — this is what allows DOS 5.0 to be loaded high, saving more memory below 640Kb. Again, this can't be done in an XT, since the 8088 processor can't address any extended memory at all.

There are a couple of things you can still do to maximise the free memory in your XT, but you won't get anywhere near the amount that you can get on your '386. For starters, do you really need DOS 4.01? DOS 3.3 will work just as well with a 20Mb hard disk, and save a few Kb of RAM at the same time. Also, consider only loading drivers and TSRs only when they are actually needed, and

then unloading them afterwards. The Mark and Release utilities make this easy, as described in 'Tech Tips' (April, 1992).

Sticky disk addendum

Your Tech Tips column offers useful advice, but the explanation in the June issue as to why a hard disk might stick after shutdown seems erroneous. When two quite flat pieces of, say, glass or metal are placed together they cohere due to molecular bonding. That is, the molecules forming optically-flat surfaces get so close to each other that the very forces that hold molecular structures together in the first place will 'stick' such surfaces firmly together.

The allusion to a partial vacuum as the causality of the problem isn't tenable since there's no agent present to create one, merely because the disk has stopped. Given that the physics of such bonding are not to be denied, your advice to use gentle tapping is doubtless the most practical solution, although extensive

If you have a PC problem that's been bugging you, put the details on paper, send them in to *Your Computer* magazine, and we'll try to help. On the other hand, if you have any advice or hints on using hardware or software that might interest others, drop us a line and we'll pass them on.

The best *Tech Tip* published each month will earn the author a \$100 voucher, redeemable at any Rod Irving Electronics or Software Express store, or by mail order from either company.

The address in either case is Tech Tips, Your Computer, PO Box 199, Alexandria NSW 2015, or by fax to (02) 317 4615.

application of a ball-peen hammer to roughen the disk is a permanent cure.

George Lindley

Thanks, George, for clarifying that one for us. Perhaps we should have held your last suggestion over for the next April issue, but by then some readers might have forgotten about the initial problem.

More memory for a '386

I realise that expanded and extended memory have been exhausted as a subject, but I have a particular problem and I am not fully satisfied with the answers I have been able to elicit from others. My machine is a 25MHz '386DX running DOS 5.0, Windows 3.1, and a complete set of Microsoft device drivers.

The motherboard is, I believe, a Chips & Technologies P386, with an AMI BIOS, and room for up to 10Mb of on-board RAM. I also own an 8Mb RAM expansion card suitable for this machine.

My problem is this: the machine was originally purchased with 2Mb of RAM. I recently decided to increase this to help run Windows and some other memory-hungry software. The budget could stretch to a further 4Mb of memory, so I planned to put an additional 2Mb on the motherboard (giving a total of 4Mb on-board) and a further 2Mb as expanded memory on the expansion card, to use as a dedicated RAM drive.

This didn't seem too complex to me (ingenue that I am), and in the initial stages, occasioned no demur from my computer dealer either. However, after several attempts by the dealer to configure the machine this way (4 x 1Mb SIPPs on the motherboard, and several 256Kb chips on the expansion card), I was informed that I will be unable to use expanded memory on this machine until every bank on the motherboard is full.

My other concern is that I am told that motherboard RAM can only be

incremented in powers of two — I cannot have 6Mb of RAM, only one, two, four, or eight. The manual for my motherboard seems to suggest otherwise, and clearly states that it has a maximum of 10Mb, (which is not a power of two).

Tony Gilbert

A lot of people are confused about the subject of memory, which is why the subject is far from dead despite being covered at length in this and other computer journals. As a result of the PC family being the outcome of a considerable number of enhancements upon the original IBM design, there are many different ways in which PC memory configured — the ideal one for any situation depending on the hardware present, the software to be used, and the type of CPU.

Firstly, it is not necessary for memory to exist only in powers of two, but in many machines, this is the way it works out, for reasons I'll get to in a moment. I know people who use computers with 5Mb, 6Mb, and 12Mb of RAM, without any trouble.

The key requirement when upgrading the memory in any computer, is that it be wide enough for the processor in the machine. For an 8088 processor, this is eight bits, for the '286 and '386SX, it is 16 bits, and for the '386DX and all '486s, it is 32 bits. This is the bus width of the processor, and is the amount of data which can be transferred to or from the memory in a single operation. For every eight bits of data which are stored in memory, there is also a ninth parity bit (the purpose of I won't go into here), just to know that it's there.

Memory chips come in sizes of powers of two (actually, for design reasons, they increment in powers of four — every second power of two). The common sizes these days are 256K bit, 1M bit, and 4M bit, and either come as single chips

one bit wide, or as SIMMs or SIPPs (SIMMs with pins) nine bits wide. So, for your '386DX, any memory expansion has to consist of 36 single chips (32, plus four parity chips), or four SIMMs or SIPPs, or a multiple thereof.

The memory already on the motherboard is not important, unless you have to remove it to make way for higher-capacity chips — it's already there, and obviously works. Since your motherboard has a total capacity of 10Mb, with 2Mb already there, I presume it has sockets for a further eight SIPPs, each of 1Mb capacity. So, to expand your motherboard memory, you could add four of these, giving a total of 6Mb, or eight, totalling 10Mb.

Your dealer is undoubtedly correct in saying that you should fully populate the motherboard before adding extra chips to the expansion board, but this should not be confused with expanded memory. Thirty-two bit memory cards for '386 motherboards are exclusively (as far as I am aware), extended memory boards, and add extra memory space to the end of that which is already on the motherboard. If the motherboard isn't fully populated, then there will be a hole in the memory space between the end of the motherboard's memory, and the start of that on the expansion card, and you won't be able to use the latter at all.

If you really need expanded (EMS) memory for some specific purpose, then on a '386 machine, the best way to achieve this is to install it as extended memory memory, and then use a memory manager such as EMM386, or QEMM, to covert some of it to expanded memory.

So in your case, put another 4Mb on the motherboard, and leave the expansion card in its box until such time as you need more than 10Mb.

This index serves for the *Your Computer* magazines as well as *Tech Tips Digest* itself. It gives a reference by month e.g., 7/91, in a standard typeface, and a *Digest* page reference e.g., 63, in bold type. A month should be cited for every entry, but items appearing in *Your Computer* after August 1992 are not included in this issue of the *Digest* being scheduled for the next.

A

Above Disk, 7/91, **63**
accelerator boards, 6/91, **61**
accessing remote PCs, 9/92
ACE chip, 12/90, **26**
Acer VGA monitor, 8/92, **120**
active partition table, 7/91, **64**
adapter segment RAM, 4/92, **102**
AddStor, 10/92
Adobe PostScript, 10/92
ADV 286, 5/92, **105**
Agiler Mouse, 8/91, **68**
Aldus PageMaker, 9/90, **13**
aliases and directories in path
 statements, 9/92
Allways, 1/92, **90**
ALR 286 Powerflex Plus, 5/91, **54**
AMI BIOS, 8/90, **11**, 11/90, **21**, 12/90, **27**,
 7/91, **63**, 9/91, **72**, 1/92, **89**, 3/92, **99**,
 5/92, **105**, 8/92, **123**
Ami Pro, 3/92, **96**
Amiga 500, 11/92
amplification, connection of internal
 speaker to external, 9/92
Amstrad 2086, 3/92, **99**
Amstrad PC 1286, 3/92, **97**, 11/92
Amstrad PC1640, 5/91, **52**, 8/91, **68**
ANSI.SYS, 8/90, **12**, 10/90, **17**, 12/90, **26**,
 4/91, **49**, 7/91, **63**, 3/92, **99**, 4/92, **104**
answering machines, 2/91, **40**
Apple and fonts, 10/92
Apple File Exchange, 7/91, **66**
Apple LaserWriter IINT, 2/92, **93**
Apple LaserWriter IINTX, 10/92
AppleTalk, 2/92, **93**
April 91, Tip of the Month, 4/91, **49**
April 92, Tip of the Month, 4/92, **104**
ARC, 8/90, **9**
archives, 8/90, **9**
archiving programs, 8/90, **9**
Arista Electronics phone line sharer,
 2/91, **40**
Artex A256C colour scanner, 7/91, **66**
Artisoft LANtastic, 1/92, **92**
ASCII codes, 4/91, **49**
ASCII, 7/90, **8**, 1/91, **29**, 11/91, **80**
AST SVGA monitor, 11/92
Astral Picture Publisher, 9/90, **13**
AT, 8/90, **8**, **11**, 10/90, **19**,
 12/90, **24**, **27**, 1/91, **31**, **35**,
 2/91, **37**, 3/91, **43**, **46**,
 4/91, **47**, **49**, 6/91, **60**,
 7/91, **64**, **65**, **66**, 8/91, **68**,
 9/91, **72**, 10/91, **75**, 11/91, **79**,
 12/91, **88**, 1/92, **89**, **90**, 2/92, **93**,
 4/92, **102**, **103**, 11/92
AT-bus, 8/90, **11**
Atari ST, 3/91, **42**
ATI EGA Wonder card, 5/92, **109**
audio and CD ROM drives, 12/92
audio tape, use for data backup, 12/92
August 91, Tip of the Month, 8/91, **70**
Australian Coherent User Group
 (ACUG), 4/92, **101**
Australian modification of BIOSes for
 different hard drives, 5/92, **106**
auto parking, 11/92, 12/91, **85**
auto-ranging, 5/91, **53**
AutoCAD, 2/91, **37**
autoexec.bat, 12/92
Auva 386SX, 9/91, **73**
AWARD BIOS, 11/90, **21**

B

back up programs, 10/91, **77**
backing up, 2/91, **38**, 3/91, **45**, 10/92
bad clusters, 4/92, **103**, 10/92, 12/92
bad sectors, 12/91, **85**, 1/92, **89**, 10/92,
 11/92
bandwidth, 2/91, **37**
bank switching, 5/91, **54**, 4/92, **102**
Basic compilers, 4/92, **101**
Basic interpreters, 4/92, **101**
Basic programming, 10/92
Basica, 12/90, **26**
batch file for backing up, 3/91, **45**
batch file to test CPU, 12/92
batch file to test screen type, 12/92
battery, 2/91, **39**
BattleChess, 6/91, **57**
BBS software, 2/92, **95**
BBS software for Xenix, 2/92, **95**
BBS, 5/91, **53**, 2/92, **95**, 3/92, **98**
BBSes, 10/92
BIOS data area, 7/90, **7**
BIOS drive support, 3/92, **99**
BIOS drive tables, 11/90, **21**
BIOS ROM, 10/90, **19**
BIOS shadowing, 7/91, **63**
BIOS table, 11/90, **21**
BIOS upgrade, 4/92, **101**
BIOS, 11/90, **21**
Bit Blitzer 123E, 7/90, **4**
bit-mapped fonts, 10/92
Bitstream fonts, 10/92
Black Box Serial - Parallel Converter II,
 2/91, **38**
boot disk, 11/91, **79**
boot-up problems, 9/91, **73**
boot-up menu, 10/90, **20**
boot-up.lzh, 10/90, **20**
Boot.sys 1.42, 11/91, **82**
Boot.sys alternative, 7/92, **114**
Boot.sys compatibility with DR DOS,
 7/92, **114**
Boot.sys, 10/90, **20**, 9/91, **73**, 12/91, **84**,
 4/92, **103**, 12/92
booting DOS and Xenix, 7/91, **64**
booting from network, 3/91, **43**
booting, 1/92, **90**
Borland Turbo Basic, 4/92, **101**
Brooklyn Bridge, 9/92
Brother HR-15 daisy wheel, 2/91, **38**
bus mouse, 12/90, **24**
bus widths, 8/92, **123**

C

C++, 7/92, **114**
C, 4/91, **47**
C, 6/92, **110**
cables, 7/90, **6**
caching controllers, 8/91, **69**
caching, hazards of, 8/92, **118**
Calmer Utilities Thunderbyte board,
 7/92, **117**
Canon A1210, 8/92, **119**
Caps Lock, 10/92
caps to lower case, 7/90, **7**
Carets/By Hand Deluxe, 9/90, **13**
case conversion, 7/90, **7**
CAT reader, 9/90, **13**
CD ROM drive, 12/92
CD ROM, 1/92, **89**
CD ROMs for multimedia, 12/92
Central Point Copy II PC Deluxe Option
 Board, 7/91, **66**
Centronics standard, 2/92, **93**, **94**
CGA compatibility, 5/92, **109**
CGA, 3/91, **43**, 8/92, **119**
checking processor and screen types,
 12/92
Checkit, 4/92, **104**
Chips and Technologies NEAT chipset,
 7/91, **63**
Chips and Technologies P386,
 8/92, **123**
Chkdsk, 10/91, **75**, 12/91, **85**
Citizen 200GX, 3/92, **96**
CMOS, 7/91, **63**, 9/91, **72**
CMOS RAM, 1/92, **89**
CMOS settings, 11/91, **79**
Coherent Operating System, 4/92, **101**
Coherent User's Group, 4/92, **101**
Coherent Users Group, 5/91, **55**
Coherent, 5/91, **55**
colour monitor as mono, 8/92, **120**
colour printing, 3/92, **96**, 8/92, **119**
colour scanner, 7/91, **66**, 9/91, **73**
colour word processor, 5/92, **108**
command mode, 10/92
Commodore 486SX-25, 12/92
Commodore fast format, 8/90, **9**
Commodore PC 40, 4/92, **103**
Commodore RAM, 4/92, **103**
Compaq 386N, 7/91, **67**
Compaq SLT, 2/92, **93**
compatibility of AMI BIOS with IDE
 drives, 3/92, **99**
compatibility of disk cache and
 compression, 10/92
Complete Fax Portable, 5/91, **55**, 12/92
compression, 4/92, **103**
CompuPro multi-user CP/M system,
 2/91, **38**
Computer Aided Technology, 9/91, **73**
config.sys, 12/92
constant speed modems, 5/91, **53**
Consumer Affairs, 11/92
contiguous files, 5/92, **105**
control codes, 11/91, **80**
control key, 11/91, **80**
conventional memory, 5/91, **54**, 4/92,
 100, 9/92
conventional RAM, 6/91, **59**
Copy II PC, 7/91, **64**, **66**
copy protected software, 7/91, **64**
copy protection, 7/91, **64**, 6/92, **111**
 11/92,
copying protected software, 7/91, **64**
Copywrite, 7/91, **64**
Corel Draw 1.2, 11/90, **21**
Corel Draw, 9/90, **13**, **16**, 8/91, **70**, 10/92
Corel Trace, 11/90, **21**
correction, to October 1992, Tech Tip of
 the Month, 12/92
corrupted boot disk, 12/91, **84**
counting in batch files, 5/92, **107**
CPU speed, 1/92, **92**
CPUs and speed, 1/92, **92**
cross platform, 1/91, **29**, 2/91, **41**, 7/91,
 66, 7/92, **114**
customising clock display, 9/92
customising Windows, 10/92
cyclic redundancy check (CRC), 6/92,
 111

This index serves for the *Your Computer* magazines as well as *Tech Tips Digest* itself. It gives a reference by month e.g., 7/91, in a standard typeface, and a *Digest* page reference e.g., 63, in bold type. A month should be cited for every entry, but items appearing in *Your Computer* after August 1992 are not included in this issue of the *Digest* being scheduled for the next.

D

damaged hard disk, 5/91, 52
damaged disks, 6/91, 60
data acquisition, 6/92, 110
data compression, 8/90, 9, 7/92, 115, 9/92
data files, 10/92
data logging, 6/92, 110
data recovery, 3/91, 42, 6/91, 60
data terminal ready (DTR), 10/92
Datamini 286 AT, 3/91, 46
Dataplex 224, 5/91, 53
date format changing, 5/91, 55
dates, Australian format, 5/91, 55
Dayna Communications DOS Mounter, 7/91, 66
DaynaFile, 1/91, 29, 2/91, 41
DB-25, 10/90, 18
dBase and DR-DOS, 3/91, 46
dBase IV v.1.1, 3/91, 46
dBase IV, 6/92, 110
DD, 11/90, 22, 8/91, 68, 12/91, 86
December 91, Tip of the Month, 12/91, 86
December 92, Tip of the Month, 12/92
default boot sequence, 1/92, 90
default boot, 7/91, 64
delayed write caches, 8/92, 118
deleting hidden files, 5/92, 106
deleting system files, 5/92, 106
DeskJet, draft mode 12/90, 25
DeskJet, 12/90, 25
DeskJet, changing default, 12/90, 25
DeskJet, default, 12/90, 25
DeskJet, draft quality, 12/90, 25
DeskJet, letter quality, 12/90, 25
DESQview BAT files, 10/90, 17
DESQview, 10/90, 17, 20, 1/91, 28, 34, 6/91, 57, 8/91, 69, 10/91, 75, 6/92, 109, 7/92, 114
DIP, 10/90, 19
dirt on drive heads, 2/92, 95
dirt on heads 7/92, 118
disk cache, 3/91, 43, 6/91, 57, 7/91, 65, 1/92, 90, 3/92, 97, 4/92, 100, 103, 8/92, 118, 11/92
disk compression and caching, 10/92
disk compression, 9/92, 10/92
disk controllers, 8/91, 69
disk densities, 12/91, 88
disk doubling, 11/90, 22
disk drive, 10/91, 75
disk drive springs, 10/90, 17
disk drive, 10/90, 17
disk housekeeping, 10/92
Disk Manager, 8/90, 11, 4/92, 104
disk partition, 11/91, 81
disk problems and solutions, 5/92, 105
disk recovery, 6/91, 60
Diskcomp, 5/92, 106
Diskcopy, 5/92, 106
disks, 7/90, 5, 8/90, 9, 11/90, 22, 12/90, 25
display mode changes, 9/92
displaying text files, 10/91, 77
dongle, 11/92
dongles, 6/92, 111
DOS + Xenix, 4/91, 47
DOS 3.2, 7/91, 63, 4/92, 104
DOS 3.3, 11/90, 21, 4/91, 47, 49, 10/91, 75, 11/91, 81, 12/91, 84, 4/92, 103, 5/92, 105, 106, 107

DOS 4 installation, 1/91, 28
DOS 4.0, 9/90, 15, 11/90, 21, 1/91, 31, 9/91, 73
DOS 4.01, 12/90, 27, 1/91, 28, 10/91, 77, 2/92, 94, 3/92, 96, 97, 5/92, 105, 8/92, 121
DOS 5.0 & parallel printers, 2/92, 94
DOS 5.0, 1/91, 33
DOS 5.0 manual, 4/92, 104
DOS 5.0 upgrade disk, 3/92, 96
DOS 5.0, 1/91, 33, 6/91, 57, 9/91, 71, 10/91, 77, 11/91, 79, 1/92, 90, 2/92, 94, 3/92, 96, 97, 4/92, 101, 103, 5/92, 106, 107, 108, 7/92, 116, 8/92, 121, 9/92, 10/92
DOS copy command, 10/91, 77
DOS Fdisk, 11/90, 21
DOS, 12/91, 85, 2/92, 95
dot-matrix printer, 3/92, 99
double density disks, 10/92
double density, (2D) or (DD), 7/90, 5
downloaded fonts, 10/92
DR DOS 6.0, 7/92, 114, 9/92, 10/92
DR DOS, 1/92, 89
DR-DOS 3.3, 3/91, 46
Dr. Genius, 6/91, 57
drive parameters, 3/92, 96
Drive Swapper, 10/92
drives, 9/90, 16
DSHD, 11/90, 22
DTK BIOS v.4.26, 4/92, 101
DTP print enhancement, 3/92, 99
DTP, 11/92, 54
dual scan, 7/90, 4
DX.COM, 9/92

E

Easycall, 8/90, 8
editing TIFF files, 7/90, 5
EGA, 12/90, 26, 1/91, 34, 8/92, 119
Eight as in:-
8-bit cards, 6/91, 61
8-bit EMS board, 7/91, 65
8-bit, 8/92, 123, 9/92
EISA bus, 9/92
EIZO 8060S, 1/91, 34
EMM386, 8/92, 123
EMM386.exe, 12/92
EMM386.SYS, 7/91, 65
EMS 3.2, 12/91, 87
EMS 4.0, 11/91, 79, 12/91, 87
EMS board, 3/91, 43
EMS, 11/90, 23, 2/91, 37, 5/91, 54, 6/91, 61, 7/91, 63, 65, 9/91, 71, 4/92, 102
enhancement of DTP printing, 3/92, 99
environment space, 7/91, 67
EPROM programming of BIOS drive table, 3/92, 99
Epson AX2, 11/90, 21
Epson EX8000, 12/90, 26
Epson LQ series, 3/92, 96
Epson LQ500, 3/92, 99
Epson LX-80, 3/92, 96
Epson PCe XT, 3/92, 97
Epson RX-100, 3/92, 99
Epson-emulation, 3/92, 96
errors, 1/91, 31
esc.com, 7/91, 63
escape character, 10/92
ESDI drives, 8/90, 11

Ethernet, 2/92, 93
Everex AT, 8/90, 8, 12/90, 27
Excel 3, 2/92, 93
Excel, 4/92, 103
executable Basic files, 4/92, 101
expanded memory board, 11/92
expanded memory, 5/91, 54, 6/91, 59, 4/92, 100, 8/92, 123, 9/92
expansion boards, 5/92, 105
extended IBM character set support, 11/92
extended keyboard, 12/90, 26
extended memory loss, 5/91, 54
extended memory, 5/91, 54, 7/91, 63, 9/91, 71, 11/91, 79, 1/92, 90, 3/92, 97, 4/92, 100, 103, 8/92, 123, 9/92
external modem, 1/91, 32

F

F1 key, 1/92, 90
F11, 12/90, 26, 4/91, 49, 7/91, 63
F11F12-.COM, 4/91, 49
F12, 12/90, 26, 4/91, 49, 7/91, 63
fast format, 8/90, 9
fast format for DOS floppies, 8/90, 9
Fastmenu Gold, 9/92
FAT, 8/90, 9, 12/90, 25, 5/91, 52, 10/91, 75, 12/91, 85
fax cards, 12/92
fax modem, 12/90, 24, 1/91, 32
fax, 2/91, 40
FaxScan, 9/90, 13
FCBs, 1/91, 28
FDHD, 2/91, 41, 7/91, 66
Fdisk, 11/91, 81
Fidonet, 2/92, 95
file control blocks (FCBs), 1/91, 28
File Manager, 1/91, 32, 3/91, 42, 5/92, 108, 7/92, 116
file handling, 1/91, 28
file security, 1/91, 31
file transfer, 1/91, 29, 7/91, 66
file translation, 7/91, 66
file translators, 1/91, 29, 2/91, 41
filename, 8/90, 11
files, 2/91, 39
financial aspects of a BBS, 3/92, 98
Finder, 7/90, 6
Five as in:-
586, 1/92, 92
Flight Simulator 3, 2/91, 39
floppies and mouses, 8/91, 68
floppies, 8/90, 9, 11/90, 22, 12/90, 25, 3/91, 42, 9/92
floppy disk capacity, 3/91, 42
floppy disk drives, 8/91, 68
floppy disk formatting, 3/91, 42
floppy disk, 3/91, 42, 2/92, 95
floppy drive errors, 10/92
floppy drive, 10/92
floppy drives, 6/92, 110
floppy drive, inoperative, 6/92, 110
font cartridge, 10/92
font definition, 10/92
font family, 10/92
fonts, 10/92
format, 12/90, 25
formatting floppies, 1/91, 31, 10/92, 8/91, 68
formatting hard disks, 8/91, 68
formatting, 8/90, 9

This index serves for the *Your Computer* magazines as well as *Tech Tips Digest* itself. It gives a reference by month e.g., 7/91, in a standard typeface, and a *Digest* page reference e.g., 63, in bold type. A month should be cited for every entry, but items appearing in *Your Computer* after August 1992 are not included in this issue of the *Digest* being scheduled for the next.

Four as in:-

486, 11/91, **79**, 1/92, **89**, **92**, 2/92, **93**,
4/92, **100**, 7/92, **118**
486DX, 11/91, **79**, 7/92, **115**, 10/92,
11/92, 12/92
486SX-25, 12/92
4DOS, 11/92
fragmentation, 10/92
fragmented files, 5/92, **105**
Freelance, 2/92, **93**
function keys, 7/91, **63**
games, 7/90, **4**, 9/90, **13**

G

Geniscan GS-C105, 12/91, **87**
Genius GM-F302 mouse, 6/91, **57**
Genius GM6 mouse, 10/91, **76**
GeoWorks Pro, 11/92
GIF, 3/92, **98**
GoScript, 3/92, **99**
graphic file formats, 3/92, **98**
graphical user interface (GUI), 6/92, **109**
graphics, 3/91, **43**
graphics cards, 7/90, **4**, 9/90, **13**
graphics mode, 6/91, **57**
Gray-FX, 9/90, **13**
GT-20 Mini 10MHz motherboard, 5/92, **105**
Guardian, 1/92, **90**
GUI, on XT, 11/92
GUI, 6/92, **109**
GW Basic, 12/90, **26**, 2/92, **94**, 4/92, **101**

H

H-P DeskJet Plus, 2/91, **39**
H-P DeskJet, 12/90, **25**, 2/92, **94**
H-P LaserJet emulation, 3/92, **96**
H-P LaserJet III, 2/91, **38**
H-P LaserJet, 2/92, **93**
H-P Printer Control Language (PCL),
10/92
H-P serial plotter, 2/92, **93**
handshaking, 7/90, **4**, 10/92, 2/92, **93**, **94**
hangup command, 10/92
hard disk bearings, 10/91, **75**
hard disk boot, 1/91, **29**
hard disk boot sector, 9/91, **73**
hard disk compressors, 10/92
hard disk connector 2/91, **37**
hard disk controller card, 12/90, **27**
hard disk drives, 10/91, **77**
hard disk maintenance, 9/90, **16**
hard disk mounting 6/91, **60**
hard disk operation, 11/92
hard disk partition, 4/91, **47**
hard disk ROM, 4/92, **102**
hard disk storage, 1/92, **92**
hard disk, 8/90, **11**, 2/91, **37**, 5/91, **52**,
6/91, **60**, 1/92, **89**, 3/92, **96**, 4/92, **103**,
8/92, **122**, 9/92, 10/92
hard drive orientation, 6/91, **60**
hard drive, 11/92, 2/91, **37**, 6/92, **110**
hardware cache, 8/91, **69**
hardware locks, 6/92, **111**
hardware virus protection, 7/92, **117**
Hayes, 10/92
HD, 11/90, **22**, 8/91, **68**, 12/91, **86**

head alignment, 1/92, **89**, 9/92
head cleaning, 7/92, **118**
head parking, 9/90, **16**, 12/91, **85**, 11/92
Hercules, 3/91, **43**
Hercules-compatible, 7/90, **4**
HiCom DriveCache, 8/91, **69**
high density (HD), 7/90, **5**
high density disks, 10/92
high level format, 8/91, **68**
high memory area (HMA), 6/91, **57**,
9/91, **71**
high memory, 6/91, **57**
high RAM, 11/91, **79**, 4/92, **103**
high resolution graphics, 11/92
Hijaak, 9/90, **13**
himem.sys, 12/92
HMA, 9/91, **71**
humour, 7/90, **6**, 8/90, **11**, 8/92, **122**

I

I/O map, 7/90, **7**
IBM 8514/A, 2/91, **37**
IBM Advanced Diagnostics disk, 4/92, **104**
IBM Professional Graphics Controller
(PGC), 5/92, **105**
IBM PVGA card, 5/92, **105**
IBM Storyboard, 3/92, **98**
Icon Draw, 11/92
icon names on the Mac, 7/90, **6**
icons, 7/90, **6**, 1/91, **32**, 11/92
IDE drive, adding a second, 7/92, **115**
IDE drive, 8/91, **69**
IDE Drives, 8/90, **11**
IDE, 8/90, **11**, 11/90, **21**, 6/91, **60**, 8/91,
6/91, 10/91, **77**, 1/92, **89**, 3/92, **99**, 7/92,
11/92, 9/92, 10/92
Image Scan/Paint, 9/90, **13**
Image-Read, 9/90, **13**
Imagineering Ultra 286, 1/92, **90**
indexing images, 6/91, **62**
installing boards, 3/91, **46**
installing chips, 3/91, **46**
installing SIMMs, 3/91, **46**
integrated drive electronics, 8/90, **11**,
11/90, **21**, 8/91, **69**
interactive batch files, 8/92, **120**, 9/92
interlaced video, 2/91, **37**
interlacing, 2/91, **37**
intermittent faults, 11/92
intermittent keyboard problems, 11/92
internal modem, 1/91, **32**
Internal or external modem, 1/91, **32**
interrupts, 4/91, **49**
inverting floppies, 12/91, **86**
InWord, 1/92, **90**
IR, 9/90, **16**
ISA, 9/92
ISys 2.0, 6/91, **62**
ISYS, 11/90, **23**

J

January 92, Tip of the Month, 1/92, **91**
July 92, Tip of the Month, 7/92, **116**
June 91, Tip of the Month, 6/91, **62**
Jun 92, Tip of the Month, 6/92, **112**
JX 10Mb hard disk, 4/92, **104**

K

key codes for F11 & F12, 12/90, **26**
key codes, 12/90, **26**
key disks, 11/92
keyboard assignments, 4/92, **104**
keyboard customisation, 8/90, **12**,
10/90, **17**, 12/90, **26**
keyboard flags, 12/92
keyboard interface errors, 11/92
Keyfake, 4/91, **47**

L

Landmark PC Probe, 4/92, **104**
Landmark speeds, 1/92, **92**, 11/92
LANtastic, 1/92, **92**
LapLink III, 8/90, **8**, 3/92, **97**
LapLink Pro, 9/92
laptop battery capacity, 2/91, **39**
laptop scanner, 12/90, **26**, 7/91, **66**
laptop, 2/91, **39**, 9/91, **73**
large partitions, 9/90, **15**
LaserWriter and Windows 2/92, **93**
letterheads, 11/91, **82**
LHARC, 8/90, **9**
Lightning 386-25, 12/92
LightScan 400J scanner, 12/91, **84**
LIM EMS, 9/91, **71**, 4/92, **102**
Linotronic, 10/92
lock program correction, 12/92
locking the hard disk, 7/92, **117**
lost clusters, 10/91, **75**, 12/91, **85**
Lotus 1-2-3 v.3.0, 9/91, **73**, 12/92
Lotus 1-2-3 v.3.1, 7/91, **67**
Lotus 1-2-3, 5/91, **54**, 6/91, **57**, **61**, 9/91,
7/91, **73**, 10/91, **76**, 11/91, **80**, 1/92, **90**
Lotus Printgraph, 11/91, **80**, 3/92, **98**
Lotus-Windows conflict, 9/91, **73**
low-end networks, 1/92, **92**
low-level format, 8/91, **68**, 10/92
low-level formatting, **39**
lower case, 7/90, **7**

M

Mac disks in a PC, 7/91, **66**
Mac file names, 7/90, **6**
Mac modem cables 7/90, **4**
Mac Plus, 7/90, **4**, 1/91, **29**
Mac to RS-232C cables, wiring diagrams,
7/90, **4**
Mace, 3/91, **42**
machine code, 4/92, **101**
Macintosh, 7/90, **4**, **6**, 1/91, **29**, 2/91, **41**,
3/91, **42**, 7/91, **66**, 10/92
MacLink Plus/PC, 2/91, **41**
macro, 1/92, **91**
MacWrite, 2/91, **41**
Maestro ZXR 2400, 10/92
Maestro, 12/91, **87**
Manifest, 1/92, **90**, 8/92, **121**
Mapmem, 4/92, **100**
March 91, Tip of the Month, 3/91, **45**
March 92, Tip of the Month, 3/92, **97**
Mark and Release, 11/90, **23**, 8/92, **121**
mass formatting of floppies, 5/92, **105**
master/slave IDE drives, 7/92, **115**
May 91, Tip of the Month, 5/91, **56**
May 92, Tip of the Month, 5/92, **107**
MCA, 9/92

This index serves for the *Your Computer* magazines as well as *Tech Tips Digest* itself. It gives a reference by month e.g., 7/91, in a standard typeface, and a *Digest* page reference e.g., 63, in bold type. A month should be cited for every entry, but items appearing in *Your Computer* after August 1992 are not included in this issue of the *Digest* being scheduled for the next.

MDA-compatible, 7/90, 4
memory, 10/90, 19, 11/90, 22, 5/91, 54, 6/91, 57, 7/91, 63, 4/92, 100
memory allocation, 11/90, 23
memory cards, 6/91, 61
memory effect, 2/91, 39
memory for Lotus, 12/92
memory management, 6/92, 109, 8/92, 121, 123, 12/92
memory manager, 9/92
memory organisation, 4/92, 102
MF2HD, 11/90, 22
MFM, 10/91, 77
Microbee CIAB, 11/92
Microbee to Mac 1/91, 29
Microbee, 1/91, 29
MicroChannel Architecture, 9/92
Micrografx Designer, 9/90, 13
Microgram combination controller board, 5/92, 106
Microsoft Driver Library, 2/92, 93
Microsoft Flight Simulator, 11/92
Microsoft Mouse driver, 4/91, 47, 10/91, 76
Microsoft mouse emulation, 8/91, 68
Microsoft mouse, 7/91, 67, 10/91, 76
Microsoft QuickBasic, 4/92, 101
Microsoft Support on Site disc, 10/92
Microsoft TrueType, 10/92
Microsoft Visual Basic, 10/92
MIDI, 7/92, 114
Mighty Cat, 8/91, 68
minimum RAM for Windows, 11/92
Miniscribe 3650, 6/91, 60
Miniscribe, 12/90, 27
Mirror, 6/92, 110
missing utility, 7/91, 63
MNP, 5/91, 53
MNP-4, 8/90, 9
MNP-5, 8/90, 9
modem cable, 7/90, 4
modem, 7/90, 4, 8/90, 8, 12/90, 24, 1/91, 32, 2/91, 40, 4/91, 49, 5/91, 53, 10/92
modem, conflict with Telecom T200 Touchphone, 8/90, 8
modifying controller card, 7/92, 115
monitor loses colour, 8/92, 120
monitor upgrade, 8/92, 119
monitors, 7/90, 4, 9/90, 13, 1/91, 34, 2/91, 37, 8/92, 119, 120, 9/92
monitors, leaving permanently on, 8/92, 120
monochrome card, 7/90, 4
more COM ports, 4/91, 49
more memory for 386, 8/92, 123
more memory for XT, 8/92, 121
motherboard memory, 2/91, 37
motherboards, 6/91, 61
mouse choices, 12/90, 24
mouse driver, 10/91, 76, 4/92, 100
mouse support for your programs, 6/92, 112
Mouse Systems driver, 8/90, 8
Mouse Systems mode, 8/91, 68
Mouse Systems optical mouse, 10/91, 76
Mouse Systems PC Paint, 3/92, 98
mouse, 12/90, 24, 26, 4/91, 47, 49, 6/91, 57, 7/91, 67, 10/91, 76, 12/91, 85
mouse/driver compatibility, 12/91, 85
MPC software, 12/92
MS Word 5.0, 10/90, 19, 12/92
MS Word 5.5, 7/91, 66

MS Word for Mac 4.0, 1/91, 29
MS Word, 2/91, 41
MS Works 2.0, 2/91, 39
MS-DOS 4.0, 4/91, 49, 6/91, 57
MS-DOS 4.01 EMS driver, 7/91, 65
MS-DOS 4.1, 7/91, 63
MS-DOS 5.0, 12/92
MS-DOS Executive, 3/91, 42
MS-DOS 5.0, 9/92, 10/92
MS-DOS, 8/90, 9, 12/90, 25, 3/91, 46, 7/91, 63, 7/92, 114, 11/92
ms.dos.exe, 3/91, 42
multi I/O controller board, 11/92
MultiFinder, 7/90, 6
multimedia, 9/92
multitasking, 6/91, 57, 6/92, 109
music programming, 7/92, 114

N

NEAT chipset, 7/91, 63, 11/91, 79, 81, 1/92, 90, 4/92, 103, 8/92, 121
NEC dot matrix, 2/92, 93
NEC MultiSync 3D, 2/91, 37
NEC MultiSync 4D, 2/91, 37
NEC MultiSync monitors, 8/92, 119
NEC P5XL, 5/92, 108
NEC Powermate 286, 9/91, 71
NEC Powermate 386/20, 4/91, 47
NetComm 1234SA, 5/91, 53
Netvet, 2/92, 95
new Boot.sys, 11/91, 82
NeXT, 2/91, 41
NiCad batteries, 2/91, 39
no-name disks, 7/92, 118
noisy fan bearings, 11/92
noisy monitor, 9/92
non-bootable DOS 5.0 disk, 3/92, 96
non-contiguous, 5/92, 105
Northman MV-4D non-interlaced SVGA monitor, 9/92
Norton Antivirus, 10/92
Norton Calibrate, 10/92
Norton Commander, 1/91, 31, 7/91, 65, 10/91, 75, 10/92
Norton Desktop, 12/92
Norton Disk Doctor, 2/92, 95, 6/91, 60, 9/91, 73, 9/92, 10/92
Norton Disk Editor, 4/91, 47, 5/91, 52
Norton Disk Tools, 4/91, 47
Norton Safe Format, 8/91, 68
Norton Speed Disk, 5/91, 52, 5/92, 105, 11/92
Norton SysInfo, 7/90, 7
Norton Utilities 4.0, 5/91, 52
Norton Utilities 4.5, 5/91, 52
Norton Utilities 5.0, 5/91, 52, 6/91, 60
Norton Utilities 6.0, 1/92, 90
Norton Utilities System Info 5.0, 4/92, 103
Norton Utilities, 7/90, 7, 3/91, 42, 4/91, 47, 5/91, 52, 5/92, 105, 10/92
Notepad, 10/91, 76, 2/92, 93
November 91, Tip of the Month, 11/91, 80
November 92, Tip of the Month, 11/92
Novell network, 2/91, 38
null modem, 10/90, 18, 1/91, 29, 2/91, 41, 4/91, 49, 2/92, 93, 9/92
Num Lock, 10/92
numeric keypad, 7/90, 8

O

Oak VGA card, 12/91, 87
October 91, Tip of the Month, 10/91, 77
October 92, Tip of the Month, 10/92
OkEditor, 5/92, 108
Olivetti PCS286, 9/92
Olivetti Typrinter, 2/92, 94
On Track Disk Manager, 11/90, 21
one wait state, 9/91, 72
Optiks, 9/90, 13
Optune Combo, 12/92
Opus, 2/92, 95
OS/2 2.0, 5/92, 105, 6/92, 109
OS/2, 4/91, 47, 8/91, 69, 12/91, 87, 11/92

P

packet radio controller, 4/91, 49
PageMaker, 7/90, 5, 7/90, 7, 2/91, 41
PaintBrush, 5/91, 54, 6/91, 62, 11/92
PaintShow, 9/90, 13
parallel printer, 2/92, 93
parallel to serial, 2/91, 38
parallel, 7/90, 7, 2/91, 38
parity errors, 1/91, 31
parity, 1/91, 31
particular user sub-directories, 10/92
partition table editor, 5/91, 52
partition table, 11/91, 81
partitioning hard disk, 5/91, 52, 7/91, 64
partitioning large drives, 11/90, 21
Pascal, 6/92, 110
password protection, 7/92, 116
pause time, 10/92
PC Draft CAD, 9/92
PC Paintbrush IV, 9/90, 13
PC Paintbrush, 11/91, 82
PC Shell, 10/91, 77
PC Tools 4.3, 5/92, 106
PC Tools, 10/92, 12/92
PC-Alien, 1/91, 29
PCX, 3/92, 98
PEM chipset, 4/92, 101
Perstor controller, 12/90, 27, 6/91, 60
Philips AT, 5/92, 105
Phoenix BIOS, 8/90, 11, 11/90, 21, 3/92, 99
phone line sharer, 2/91, 40
Pic, 3/92, 98
Picem, 6/91, 62
PKARC, 8/90, 9
PKLite, 1/92, 92
PKPAC, 8/90, 9
PKUNPAK, 8/90, 9
PKUNZIP, 8/90, 9
PKWare, 8/90, 9
PKXARC, 8/90, 9
PKZIP, 8/90, 9, 9/92
PostScript, 2/92, 93
PowerMenu, 7/92, 114
President AT, 1/91, 31
Priam 803-11 SCSI drive, 2/91, 37
print buffer, 1/92, 92
print cache, 6/91, 57
print capture, 12/90, 26
print spooler, 3/91, 43, 5/91, 54, 11/92
print spooling, 1/92, 92
PrintAPlot, 12/90, 26
PrintCache, 11/92, 54
printer, 7/90, 7

This index serves for the *Your Computer* magazines as well as *Tech Tips Digest* itself. It gives a reference by month e.g., 7/91, in a standard typeface, and a *Digest* page reference e.g., 63, in bold type. A month should be cited for every entry, but items appearing in *Your Computer* after August 1992 are not included in this issue of the *Digest* being scheduled for the next.

printer compatibility, 11/92
printer control codes, 11/91, **80**
printer driver problem, 2/91, **39**
printer driver, 2/91, **39**
printer emulation, 3/92, **96**
printer information request 8/92, **119**
printer initialisation, 1/92, **89**
printer resolution, 10/92
printer resident fonts, 10/92
printers, 7/90, **7**, 12/90, **25**, **26**, 2/91, **38**,
 3/91, **46**, 6/91, **57**, 11/91, **80**, **82**,
 1/92, **89**, 2/92, **93**, **94**, 3/92, **96**, **99**,
 5/92, **108**, 10/92
printers and fonts, 10/92
printing graphics characters, 3/92, **99**
print spooler, computer as a, 1/92, **92**
Procomm 2.4.2, 10/92
Program Manager, 1/91, **32**, 3/91, **42**,
 4/91, **108**, 7/92, **116**, 9/92, 11/92
Protech 486, 11/91, **79**
Publish It! Lite, 3/92, **97**
Publish It!, 3/92, **99**
Publishers' Powerpack, 3/92, **99**
Putkey, 4/91, **47**

Q

QBasic, 2/92, **94**
QDOS III, 12/92
QEMM 5.1, 6/91, **57**
QEMM crashes, 1/91, **34**
QRAM & extended memory, 9/91, **71**,
 1/92, **90**
QRAM, 9/91, **71**, 11/91, **79**, 1/92, **90**
Quantum IDE, 5/92, **106**
Quarterdeck Expanded Memory
 Manager, see **QEMM**
QEMM, 10/90, **20**, 1/91, **34**, 6/91, **57**,
 7/91, **65**, 8/92, **121**, 8/92, **123**
Quarterdeck QRAM, 9/91, **71**
Quarterdeck, 6/92, **109**
Quickset, 7/91, **63**
Quicksilver, 6/92, **110**

R

RAM cache, 8/91, **70**, 9/91, **72**
RAM chip speed, 5/92, **105**
RAM chips, 1/91, **31**, 3/91, **46**, 4/91, **47**
RAM cram, 3/92, **97**
RAM disk and Windows, 3/91, **43**
RAM disk, 11/90, **22**, 3/91, **43**, 6/91, **57**,
 8/91, **70**, 1/92, **90**, 3/92, **97**, 4/92, **103**
RAM drive, 4/92, **100**, 8/92, **123**
RAM needed for OS/2, 6/92, **109**
RAM, 3/91, **43**, 5/91, **54**, 6/91, **57**, 7/91,
 6/92, **97**, 4/92, **102**
ramdrive.sys, 5/91, **54**
read only cache, 8/92, **118**
read/write cache, 8/92, **118**
reading 1.44Mb disks, 7/90, **5**
reading a bent disk, 3/91, **42**
Readright OCR, 9/90, **13**
real world speeds, 11/92
Red Ryder, 1/91, **29**
reel-to-reel tape recorder as backup
 device, 12/92
reformatting hard disks, 9/90, **16**
reformatting, 1/92, **89**, 10/92
relative speeds, 11/92
remote install in LapLink III, 3/92, **97**

removable hard disks, 7/92, **117**
ResEdit, 7/90, **6**
resource editing, 7/90, **6**
Reversi, 11/90, **22**
RGB monitor for EGA, 5/92, **109**
RISC, 1/92, **92**
RJ-11, 10/90, **18**
RLL controller card, 9/92
RLL disk controllers 9/92
RLL, 10/91, **77**
Roland MT32, 7/92, **114**
Roland Sound Canvas, 7/92, **114**
Roland U220, 7/92, **114**
ROM disk, 1/92, **90**
ROM, 4/92, **102**
root directory, 10/92
RS 232, 2/92, **93**, 7/90, **4**
RS 232c interface, 6/92, **110**
RS 422, 7/90, **4**
rulers in Corel Draw, 9/90, **16**
run time Windows, 7/92, **116**

S

S-Registers, 10/92
Samsung 286, 11/91, **79**
Samsung SPC 6100 286, 9/91, **71**
scalable fonts on non-PostScript
 printers, 3/92, **99**
scalable fonts, 10/92
Scan 86, 10/92
scanner compatibility, 12/91, **87**
scanner, 12/91, **87**
SCAT chipset, 1/92, **90**
screen codes, 11/91, **80**
screen customisation, 8/90, **12**, 10/90,
 1/92, 12/90, **26**
SCSI card, 1/92, **89**
SCSI drives, 8/90, **11**
SCSI, 12/92
SEA, 8/90, **9**
Seagate 16-bit RLL controller, 9/92
Seagate ST-125A, 11/92
Seagate ST151, 7/92, **117**
Secom key, 11/92
second COM port, 12/90, **26**
second floppy controllers, 5/92, **106**
security, 1/91, **35**
September 92, Tip of the Month, 9/92
serial mouse, 12/90, **24**
serial ports, 8/90, **8**, 12/90, **26**
serial, 8/90, **8**, 2/91, **38**, 4/91, **49**
shadow RAM, 10/90, **19**, 7/91, **63**, 9/91,
 7/92, 11/91, **79**, 4/92, **100**
shadow ROM, 10/90, **19**
Share & DOS 4.0 1/91, **28**
Share, 1/91, **28**
sharing floppies, 9/92
SideKick, 11/90, **23**
SIMMs, 10/90, **19**, 8/92, **123**
Simplex, 2/92, **95**
SIPPs, 8/92, **123**
Sixteen as in:-
 16-bit bus bottleneck, 9/92,
 16-bit cards, 6/91, **61**
 16-bit, 8/92, **123**, 9/92,
 16/32 bus, 9/92,
SLEd, 9/90, **13**
Slim, 10/92
SmartDrive, 7/91, **63**, 8/91, **69**, 8/92,
 11/92, 10/92

smartdrive.sys, 3/91, **43**, 5/91, **54**, 7/91,
 6/92, 8/91, **70**
SoftBytes, 12/91, **87**
software installation, 2/91, **39**
software virus protection, 7/92, **117**
Sony KX-14CP1 RGB monitor, 5/92, **109**
Sound Galaxy NX sound card, 12/92
SoundBlaster, 12/91, **87**
SpeedStor, 8/90, **11**
Spinrite, 1/92, **89**, 3/92, **96**
spooling printing, 12/92
Spot OCR, 9/90, **13**
Sshrink, 4/92, **103**
Sxexpand, 4/92, **103**
Stacker, 10/92
standard forms with Ventura, 11/91, **80**
standard mode, 9/92
Star NX-1000, 3/91, **46**, 11/91, **80**
starting a BBS, 3/92, **98**
startup screen disabling, 3/91, **43**
startup screen, 3/91, **43**
static electricity, 3/91, **46**
static protection, 3/91, **46**
stepper motor drive, 6/91, **60**, 9/92
stepper motor hard drives, 1/92, **89**
stepper motors, 11/92, 9/90, **16**
stiction, 6/92, **110**
stuck hard disk, 6/92, **110**
styles, 10/92
Sun, 2/91, **41**
Super EGA, 8/92, **119**
Super PC-Kwik, 10/92
Supercom, 6/92, **110**
SuperStor, **38**
SuperStor 2.0, 1/92, **92**, 4/92, **103**, 9/92,
 10/92
Super VGA, 12/91, **87**
support for extended-ASCII character
 printing, 3/92, **99**
swap file, 9/90, **15**
Syncomp 386, 10/91, **75**
Syncomp Mega 386i, 8/91, **69**
Sysedit, 10/91, **76**
sysedit.exe, 3/91, **42**
Sysgen controller, 5/92, **106**
system BIOS, 9/91, **72**
system RAM, 6/91, **57**
system requirements for DTP or CAD,
 7/92, **115**
system requirements for high school
 student, 7/92, **115**
system ROM, 4/92, **102**

T

T1000SE, 1/92, **90**
Tagged Image File Format, 7/90, **5**,
 9/90, **13**, 11/90, **21**
Tandon TM-262 hard disk, 10/91, **75**
Tandy DMP 400, 2/92, **94**
Tandy Smartwatch, 11/92
tape backup, 12/92
tape drives, 2/91, **38**
Teac 1.2Mb floppy drives, 12/91, **88**
Teac 1.44Mb drive, 7/91, **66**
Teac HD drive, 8/91, **68**
Tecmar, 2/91, **37**
Telcom, 1/91, **29**
telephone line sharing, 2/91, **40**
Telix 3.12, 8/90, **8**
Telix, 4/91, **49**, 5/91, **53**, 10/91, **75**
temporary swap file, 9/90, **15**

This index serves for the *Your Computer* magazines as well as *Tech Tips Digest* itself. It gives a reference by month e.g., 7/91, in a standard typeface, and a *Digest* page reference e.g., 63, in bold type. A month should be cited for every entry, but items appearing in *Your Computer* after August 1992 are not included in this issue of the *Digest* being scheduled for the next.

Terran 486, 2/92, **93**
text files, 10/91, **77**
text mode, 6/91, **57**
Thirty as in:-
32-bit data path, 6/91, **61**
32-bit, 8/92, **123**, 9/92
Three as in:-
386 Enhanced Mode, 9/92
386 memory usage, 6/91, **57**
386, 7/90, **7**, 9/90, **15**, 10/90, **17**, 11/90, **21**, **23**, 12/90, **26**, 1/91, **28**, 3/91, **43**, 4/91, **49**, 5/91, **54**, 6/91, **57**, 6/91, **61**, 7/91, **67**, 9/91, **71**, 11/91, **79**, 1/92, **89**, **92**, 3/92, **96**, **97**, **99**, 5/92, **108**, 8/92, **121**, 9/92,
386DX AMD, 11/92
386DX, 10/90, **19**, 6/91, **59**, 6/91, **61**, 7/91, **65**, 8/91, **68**, **69**, 9/91, **73**, 12/91, **88**, 4/92, **101**, 5/92, **106**, 7/92, **114**, **115**, **117**, 8/92, **123**, 9/92, 10/92, 11/92, 12/92
TI MicroLaser with PostScript, 2/92, **93**
TIFF specifications, 9/90, **13**
TIFF, 7/90, **5**, **7**, 9/90, **13**, 11/90, **21**, 5/91, **54**, 3/92, **98**
Toshiba desktop docking adapter, 7/91, **66**, 9/91, **73**
Toshiba dot-matrix, 2/91, **38**
Toshiba T1000SE, 3/91, **45**, 1/92, **90**
Toshiba T1200, 1/91, **32**, 2/91, **39**
Toshiba T1200XE, 4/91, **49**, 12/92
Toshiba T1600, 7/91, **66**, 9/91, **73**
Travelling Software LapLink III, 2/91, **41**
Trident 8900C VGA card, 11/92
Trident TVGA Super VGA card, 2/91, **37**
Truename, 5/92, **105**
Tseng Labs ET4000 VGA card, 4/92, **101**
Tseng Labs ET4000 SVGA card, 9/92
Tseng Labs VGA card, 12/91, **87**
TSRComxx.ZIP, 4/92, **100**
TSRs & memory allocation, 11/90, **23**
TSRs, 11/90, **23**, 4/92, **103**, 8/92, **121**
TTL, 7/90, **4**, 9/90, **13**, 1/91, **34**
turbo board compatibility, 5/92, **105**
Turbo Pascal for Windows, 12/91, **84**
turbo speed, **39**
Twenty as in:-
20/20, 4/91, **47**, 5/91, **55**
Two as in:-
286, 10/90, **19**, 1/91, **31**, 2/91, **37**, 3/91, **43**, 3/91, **46**, 5/91, **54**, 6/91, **57**, **61**, 7/91, **63**, 8/91, **68**, **69**, 9/91, **71**, **72**, 11/91, **79**, 12/91, **87**, 1/92, **90**, 1/92, **92**, 5/92, **105**, 8/92, **120**, 9/92, 10/92, 11/92
286 DOS-compatible mode, 9/91, **71**
286 emulation, 12/91, **87**
286 protected mode, 9/91, **71**, 11/91, **79**
286 real mode, 11/91, **79**
2D, 11/90, **22**
two IDE drives, 7/92, **115**
two modems, 8/90, **8**
Type 1 fonts, 10/92

U

UART, 8/90, **8**
undeletable directory, 5/92, **106**
uninterruptible power supply (UPS), 12/91, **85**
Unix, 5/91, **56**

unlocking keyboard locks, 10/92
unreadable disk, 2/92, **95**
unrecoverable application errors (UAEs), 11/92
unreliable power supply, 12/91, **85**
updated BIOS, 3/92, **99**
US format dates, 5/91, **56**
utility, 7/91, **63**

V

V.32 M5, 5/91, **53**
Ventura 3.0 for Windows, 6/91, **59**
Ventura Professional Extension, 6/91, **59**
Ventura Publisher, 7/90, **8**, 9/90, **13**, 11/91, **82**
Ventura under Windows, 6/91, **59**
Ventura Publisher v.2, 6/91, **59**
VESA, 12/91, **87**
VGA card, 5/92, **109**
VGA on old monitor, 1/91, **34**
VGA, 7/90, **4**, 12/90, **26**, 1/91, **34**, 8/92, **119**
video BIOS, 9/91, **72**
video RAM, 7/91, **67**, 4/92, **102**
Vidi PC, 12/91, **87**
Virtual Memory Manager, 9/92
virtual memory, 6/92, **109**, 9/92
virtual page, 6/91, **57**
virus, 10/91, **77**, 1/92, **89**, 7/92, **117**
Vision Mouse, 10/91, **76**
voice coil actuation, 11/92
voice coil drives, 9/90, **16**, 12/90, **27**, 6/91, **60**, 12/91, **85**
volume labels, 1/91, **31**
VP-Planner, 11/91, **80**

W

wait states, 9/91, **72**, 1/92, **92**
warranty service, 11/92
Western Digital SVGA card, 12/92
Western Digital ZPVGa card, 11/92
Windows, 7/90, **5**, 7/90, **7**, 11/90, **22**, 1/91, **28**, 3/91, **42**, 6/91, **61**, 7/91, **65**, 8/91, **69**, 8/91, **70**, 12/91, **87**, 3/92, **96**, 4/92, **100**, 10/92, 11/92
Windows 1.x, 3/91, **42**
Windows 2.0, 2/92, **93**
Windows 2.11, 5/91, **54**, 7/91, **63**
Windows 2.x, 3/91, **42**
Windows 3.0 Enhanced Mode, 4/91, **49**
Windows 3.0 in Real Mode, 11/92
Windows 3.0, 7/90, **7**, 9/90, **15**, 10/90, **19**, 1/91, **31**, **32**, **34**, 3/91, **42**, **43**, 5/91, **54**, 6/91, **57**, **59**, 7/91, **63**, **67**, 8/91, **68**, 9/91, **72**, **73**, 10/91, **76**, 1/92, **90**, 2/92, **93**, 5/92, **108**, 7/92, **116**, 8/92, **118**, 9/92, 11/92
Windows 3.1, 6/92, **109**, 8/92, **123**, 9/92, 10/92, 12/92
Windows and disk formats, 8/91, **68**
Windows conflicts with DOS drivers, 4/92, **100**
Windows Desktop, 7/92, **116**
Windows enhanced mode, 5/92, **108**, 6/92, **109**
Windows files, 12/91, **84**
Windows for an XT, 11/92
Windows icons, 1/91, **32**

Windows in Advanced Mode, 7/91, **67**, 10/91, **76**
Windows in Enhanced Mode, 6/91, **59**, 7/91, **65**, 7/91, **67**, 10/91, **76**, 9/92
Windows in Real Mode, 6/91, **59**, 8/91, **68**, 12/91, **84**, 9/92
Windows in Standard Mode, 6/91, **59**, 7/91, **65**, 9/92
Windows Notepad, 3/91, **42**
Windows Print Manager, 11/92 **54**
Windows problems, 10/91, **76**
Windows recorder, 1/92, **91**
Windows start-up time, 12/92
Windows startup, 3/91, **43**
Windows swap files, 9/90, **15**
Windows with Multimedia, 1.0, 9/92
Wonunder single-card adapter, 7/91, **66**, 9/91, **73**
Word 5.0a, 10/91, **76**
Word for Windows 1.0, 7/91, **63**
Word for Windows 1.1, 7/91, **63**
Word for Windows, 5/91, **54**, 7/91, **65**, 2/92, **95**, 3/92, **96**
Word for Word, 1/91, **29**
word processor support for colour, 3/92, **96**
WordPerfect, 2/92, **93**, 12/92
Wordstar 3.1, 1/91, **29**
Wordstar 4.0, 7/90, **7**, 11/91, **80**
WordStar 5.5, 10/91, **75**
WordStar, 1/91, **28**, 2/91, **41**, 10/91, **75**, 2/92, **94**, 3/92, **96**
World Atlas, 3/92, **97**
WorldPort, 12/90, **24**, 1/91, **32**, 3/92, **96**
write-back caches, 8/92, **118**
Wyse 386, 9/90, **15**
Wyse DOS 3.31, 9/90, **15**

X

XCMOS NEAT chipset, 9/91, **73**
Xenix, 4/91, **47**, 8/91, **69**, 2/92, **95**
XModem protocol, 1/91, **29**
XMS, 12/92, 9/91, **71**
XT memory expansion, 6/91, **61**
XT EMS RAM board, 1/92, **90**
XT memory, 4/92, **102**
XT turbo, 8/92, **121**
XT, 7/90, **4**, 9/90, **13**, 10/90, **19**, 12/90, **24**, 12/90, **26**, 1/91, **35**, 3/91, **43**, 4/91, **47**, 6/91, **61**, 7/91, **65**, 8/91, **68**, 10/91, **75**, 11/91, **79**, **80**, 1/92, **89**, **92**, 3/92, **97**, **99**, 4/92, **102**, **104**, 5/92, **105**, 6/92, **110**, 7/92, **115**, **118**, 9/92, 11/92, 12/92
XTree Gold, 7/91, **67**, 10/91, **76**, 10/92, 12/92

Z

zero wait state, 9/91, **72**
ZIP, 10/90, **19**

For the Hard to Find Components
RIE is the Answer.

Computer Systems
Computer Hardware
Computer Software

Disks & Accessories

Plus a Huge range of
Home Electrical Products,
Electronic Components
& Electrical Hardware.
Over 7000 Line Items!

DON'T FORGET OUR OAKLEIGH STORE: 240C HUNTINGDALE RD.

5.25" from 39c OVER 16,000,000 SOLD! 3.5" from 62c

BULK DISK PRICES

ALL PRICES PER BOX OF TEN - WITH LIFETIME WARRANTY

DESCRIPTION	1-9	10+	50+	100+	500+
5 1/4" DS/DD	\$4.75	\$4.60	\$4.40	\$4.20	\$3.95
5 1/4" DS/HD	\$8.50	\$8.25	\$8.10	\$7.50	\$6.90
3 1/2" DS/DD	\$7.90	\$7.20	\$6.90	\$6.50	\$6.25
3 1/2" DS/HD	\$13.50	\$12.90	\$11.95	\$10.50	\$9.90

HARD DRIVES

PRICES WITHOUT PADDLE CARDS & BRACKETS	
40M DRIVE SEAGATE.....	\$299
85M HD 6ms T.T ACCESS. VOICE COIL	
2 YEAR WARRANTY.....	\$399
105M HD 6ms T.T ACCESS.....	\$469
120 KALOCK DRIVE.....	\$449
126M HD 6ms T.T ACCESS. VOICE COIL	
2 YEAR WARRANTY.....	\$539
170M 6ms T.T 2 YEAR WARRANTY.....	\$649
212M HD 5ms T.T ACCESS. VOICE COIL	
2 YEAR WARRANTY.....	\$729
345M HD 5ms T.T ACCESS.....	\$1295

Canon
THE LITTLE SQUIRT
BJ-10ex

Only \$599
\$499 tax ex

EDUCATION SOFTWARE SPECIALS.

Word Perfect for Windows.....	\$215.00
Coral Draw V.3.....	\$340.00
MS Excel V.4.....	\$290.00
Pagemaker 4.0.....	\$365.00
dBase IV.....	\$280.00

Student ID Required when ordering

SHAREWARE SOFTWARE
A LARGE RANGE OF
SHAREWARE SOFTWARE IS
AVAILABLE AT A BECKETT
ST. NORTHCOTE, OAKLEIGH,
STANMORE & BLUESTAR
COMPUTERS
FROM ONLY \$3.95 FOR
5 1/4" DISKS! OVER 1500
TITLES AVAILABLE!

BACK TO SCHOOL SAVINGS!

212M MAXTOR HARD DRIVE only \$695	
EPSON LQ-100 24 pin Printer	\$333 exTax
200cps d.s. \$399	
EPSON LQ 570 24 pin Printer	
252 cps d.s \$549	\$469 ex Tax
GENISCAN	
Black & White Scanner with OCR software	\$249
SOUND BLASTER V2 (without speakers)....\$129.00	
SOUND BLASTER PRO...\$299	
CD ROM PACKS FROM...\$899.00	
SWITCH BOXES	
RS232 2 WAY.....	\$39
RS232 4 WAY.....	\$59
RS232 X OVER.....	\$59
RS232 2 WAY AUTO.....	\$69
CENTRONICS 2 WAY.....	\$39
CENTRONICS X-OVER.....	\$59
CENTRONICS 2 WAY AUTO.....	\$69
CENTRONICS 4 WAY.....	\$69
VIDEO SWITCH BOX	
DB9 + 5 PIN DIN.....	\$69.95
VGA KEYBOARD SWITCH BOX	
2 WAY.....	\$59.95
4 WAY.....	\$69.95
9 PIN SWITCH BOX	
2 WAY.....	\$49.95
4 WAY.....	\$59.95

SAVE BIG ON MOTHERBOARDS

286-16/21.....	\$119
386SX-25/31.....	\$190
386SX-33/41.....	\$205
386-33 64K CACHE.....	\$349
386-40 64K CACHE.....	\$369
386-33 128K CACHE.....	\$425
486SX-20.....	\$445
486-33 256K CACHE.....	\$1195
486-50 256K CACHE.....	\$1495
486-66 256K CACHE.....	\$1795
486 EISA.....	\$2995

IBM JOYSTICK

DURABLE, ELEGANT
DESIGN. AUTO
RETURN CENTERING.
only \$19.95

PARADISE CARD
Run Windows up to 15 times
faster.....\$275.00

JAPANESE DRIVES

1.44MB 3.5" F.D.D.....	\$89.00
1.2MB 5.25" F.D.D.....	\$99.00
101 KEYBOARDS.....	\$49.00

OKILASER 400

\$1.199

4 p/p/m laser
printer.
LED technology,
25 resident fonts.
Recycles its own toner! Very
low cost per copy!
(About 1.8c per copy)
5 YEAR WARRANTY ON LED
HEAD! TONER
CARTRIDGES...\$54.95

MAESTRO MODEMS

2400XR.....	\$299
2400ZX.....	\$399
*9600XR FAX/DATA.....	\$399.00
SUPER EXECUTIVE.....	\$599.00

*Available in Amiga, IBM, Mac.

SOUND BLASTER

16 ASP	
TRUE 16-BIT STEREO SOUND	
ASP (ADVANCED SIGNAL PROCESSING) - CD ROM	
INTERFACE - FREE	
MICROPHONE.....	\$599.00

FAX ROLLS

210x11.5x30m.....	\$7.95
210x25.4x50m.....	\$7.95
210x25.4x100m.....	\$17.95
216x25.4x30m.....	\$7.95
216x24.5x50m.....	\$9.95
210x30x25m.....	\$7.95

ECLIPSE

DISK BOXES	
5 1/4" x 100	
3 1/2" x 40	
only \$6.95 each	

SAVE ON COMPUTER SYSTEMS! - PLUS 4 YEAR P & L WARRANTY!

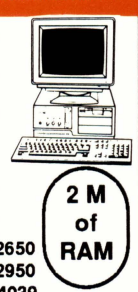
LOOK AT WHAT YOU GET!

ALL SYSTEMS INCLUDE - 2 or 4 MEG RAM
- 101 KEYBOARD - SUPER VGA (1024 x 768)
(0.28" DOT PITCH) COLOUR MONITOR -
2 SPG PORTS - 512K VGA CARD
- 1.2 MEG 5 1/4" DISK DRIVE.

DR DOS 6.0 AN EXTRA.....	\$75.00
DOS 5 AN EXTRA.....	\$99.00
DOS 5 WITH WINDOWS 3.1	
AN EXTRA.....	\$179.00

85Mb SYSTEMS

386SX-25Mhz (LM31).....	\$1399
386SX-33 (LM41).....	\$1439
386-33 (LM57) 64KCache.....	\$1519
386-33 128K Cache.....	\$1539
386-40 (LM65) 64K Cache.....	\$1539
486SX-20 (LM85).....	\$1589
486DX-33 (LM157)	
256K Cache.....	\$2279
486DX-50 (Im200+)+256KCache.....	\$2650
486DX-66 (Im200+)+256KCache.....	\$2950
486-33 EISA.....	\$4039



2 M of RAM

HAVING 16 YEARS IN THE BUSINESS, WE STAND BEHIND OUR WARRANTY.

105Mb SYSTEMS

386SX-25Mhz (LM31).....	\$1495
386SX-33 (LM41).....	\$1535
386-33 (LM57) 64K Cache.....	\$1625
386-33 128KCache.....	\$1645
386-40 (LM65) 64K Cache.....	\$1645
486SX-25 (LM85).....	\$1695
486DX-33 (LM157)	
256K Cache.....	\$2359
486DX-50 (Im200+)+256K C.....	\$2725
486DX-66 (Im200+)+256K C.....	\$3025
486-33 EISA.....	\$4229



2 M of RAM

131Mb SYSTEMS

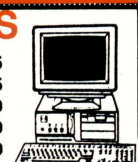
386SX-25Mhz (LM31).....	\$1550
386SX-33 (LM41).....	\$1590
386-33 (LM57) 64K Cache.....	\$1675
386-33 128KCache.....	\$1695
386-40 (LM65) 64K Cache.....	\$1695
486SX-25 (LM85).....	\$1745
486DX-33 (LM157)	
256K Cache.....	\$2435
486DX-50 (Im200+)+256K.....	\$2799
486DX-66 (Im200+)+256K C.....	\$3099
486-33 EISA.....	\$4279



4 M of RAM

200Mb SYSTEMS

386SX-25Mhz (LM31).....	\$1795
386SX-33 (LM41).....	\$1835
386-33 (LM57) 64K Cache.....	\$1920
386-33 128K Cache.....	\$1940
386-40 (LM65) 64K Cache.....	\$1950
486SX-25 (LM85).....	\$1990
486DX-33 (LM157)	
256K Cache.....	\$2680
486DX-50 (Im200+)+256K.....	\$3045
486DX-66 (Im200+)+256K C.....	\$3345
486-33 EISA.....	\$4539



4 M of RAM

Microsoft Mouse & Microsoft Windows 3.1

Package. Serial & PS/2 Version for
MS-DOS & Windows operating Systems
Includes adaptor for connection to a
PS/2-style Mouse Port. Our Price \$169.

THE FAMOUS QUICK MOUSE

Inc. Driver & test program.
- Quick Mouse Pop-Up
menu Generator.
- Quick Mouse Manual
- Image 72 Software.....\$29.95

DYNAMO & SOLAR

AM/FM RADIO WITH 4
WAY POWER BACK-UP
DOESN'T NEED BATTERIES!
WAS \$49.95 NOW JUST \$39.95
SIMM SOCKETS
With metal clips
Single.....\$5.95 Double.....\$6.95

MEMORY

MEMORY 1-9 10-99 100+	
41256-08.....	\$2.95 \$2.75 \$2.50
44256-07.....	\$7.95 \$7.50 \$6.50
SIMMS 1-9 10-24 25+ 100+	
1M x 9-60	\$62 \$60 \$59 \$57
1M x 9-70	\$59 \$57 \$55 \$54
256K-70	\$23 \$21 \$19 \$17
256K-80	\$21 \$19 \$17 \$15
4M x 9-70	\$249 \$239 \$229 \$219
SIPPS	
1M x 9-80	\$99 \$89 \$85 \$84
256K-60	\$29 \$27 \$26 \$25

RIE IS WHERE YOU FIND
9 CHIP SIMMS & SIPPS.

1M x 9-70 SIMM \$99 \$85

1M x 9-70 SIPP \$99 \$85

MATH CO-PROCESSOR

INTEL 387SX/SL.....\$189.00

INTEL 386DX.....\$199.00

16550 UART with FIFO's

For high speed modem users\$23.95

NEW FLOPPY LOCK

Keep prying eyes &
meddling hands out of
your data with a floppy
lock in 5 1/4" & 3 1/2"
.....\$39.95

"THROW OUT THOSE WELDING GLASSES!"

SVGA
MONITOR
.28" DOT PITCH
WITH A BUILT IN
ANTI-GLARE SCREEN
FILTER.....\$469.00

EPSON PRINTERS

LX-400 9 PIN.....	\$245
LX-850 9 PIN.....	\$375
LQ-580 24 PIN.....	\$895
LQ-870 24 PIN.....	\$950
LQ-1070 24 PIN.....	\$770
LQ-550 24 PIN.....	\$575
LQ-1170 24 PIN.....	\$1149

ROD IRVING ELECTRONICS Pty. Ltd

FOR THE SERIOUS COMPUTER USER! ESTABLISHED 1977.

MELBOURNE: 48 A'Beckett St. Ph: (03) 663 6151. Computers: Ph 639 1640

OAKLEIGH: 240C Huntingdale Rd. Oakleigh. Ph: (03) 562 8939

NORTHCOLE: 425 High St. Ph: (03) 489 8866

SYDNEY: 74 Paramatta Rd. Stanmore. N.S.W. Ph: (02) 519 3134 or 565 1458

ADELAIDE: 241-243 Wright St. Adelaide. Ph: (08) 211 7200

MAIL ORDER: Ph: (03) 543 7877 (LOCAL CALLS)

Mail Order Hotline: Ph: 008 33 5757 STD ORDER

BLUESTAR COMPUTERS:

MELBOURNE: 271 Maroondah Hwy. Ringwood. Ph: (03) 870 1800

SYDNEY: 115-117 Parramatta Rd. Concord. Ph: (02) 744 5526

GOVERNMENT, CORPORATE SALES & TAX EXEMPT SALES. PH: (03) 543 2166

Errors & Omissions excepted IBM® PC® XT® AT® Microsoft® are all registered trademarks of International Business Machines. PRICES & SUBJECT TO CHANGE WITHOUT NOTICE

Relax...



...with Oshima Computers, you're always a jump ahead.

OSHIMA



50486DX Super Personal Computer

- Intel 80486DX-50MHz CPU
- EISA 32-bit Local Bus Slot
- Tseng Lab ET-4000 32-bit EISA VGA Card
 - 15603 Chr/mSec speed, supports Hi-colour
 - 500% faster than standard VGA card
 - 300% faster than Paradise Accelerator
- ACERVIEW 34T Non-interlaced XGA monitor
- 4MB RAM with 256K cache
- 1.2/1.44MB floppy drive
- Super Cache IDE controller
 - Speed up your HDD at 0.3mSec seek time
- Western Digital 210MB HDD
- 2 serial, 1 parallel, 1 games port
- Honeywell industrial keyboard
- Microsoft Serial Mouse
- Microsoft DOS 5.0

Test results

	Speed (Chr/mSec)	Windows
Trident 1MB Card	2642	14 sec
Paradise Accelerator	5342	13 sec
Tseng Lab ET-4000	5342	13 sec
SPEED Tseng Lab EISA Local Bus	15603	3 sec

OSHIMA COMPUTERS Pty Ltd

Oshima Computers Pty Ltd
56-58 Burwood Road
Hawthorn 3122
Ph: (03) 819 6060
Fax: (03) 819 4124

Alpha Research Computers
Suite 1/332 Main Street
Mornington
Ph: 059 75 3577
Fax: 059 75 3902

KCT Computers
380 Victoria Street
Richmond
Ph: (03) 427 7209
Fax: (03) 428 0646

LogiMax Pty Ltd
400 Sussex Street
Sydney NSW
Ph: (02) 281 1388
Fax: (02) 281 1831